arXiv:1608.03647v2 [cs.LG] 23 Apr 2017

Learning with Value-Ramp

Tom J. Ameloot* Jan Van den Bussche

Abstract

We study a learning principle based on the intuition of forming ramps. The agent tries to follow an increasing sequence of values until the agent meets a peak of reward. The resulting Value-Ramp algorithm is natural, easy to configure, and has a robust implementation with natural numbers.

Contents

1 Introduction **2**

2 Value-Ramp algorithm **4**
2.1 Basic definitions 4
2.2 Desired properties 5
2.3 Concrete algorithm 7
2.4 Tasks formalized 9

3 Exploration on deterministic tasks **11**
3.1 Optimal value estimation 11
3.2 Optimal path following 16
3.3 Proof of Theorem 3.1 18
3.4 Proof of Corollary 3.2 22
3.5 Proof of Theorem 3.10 24

4 Greedy navigation under nondeterminism **25**
4.1 Greedy navigation 25
4.2 Reducibility 26
4.3 Restartability 28
4.4 Navigation result 30
4.5 Proof of Theorem 4.6 31

5 Conclusion and further work **38**

Appendix **42**

*This work was made while T.J. Ameloot was a postdoctoral fellow of the Research Foundation – Flanders (FWO), at Hasselt University and the transnational University of Limburg (Belgium).

1 Introduction

In reinforcement learning, techniques such as temporal difference learning (Sutton, 1988) are used to model biological learning mechanisms (Potjans et al., 2011; Schultz, 2013, 2015). In that context, Frémaux et al. (2013) have simulated neuron-based agents acting on various tasks, where the firing frequency of some neurons represents the value of encountered states. Frémaux et al. (2013) observed that the simulated value neurons behave in a ramp-like manner: the firing frequency of a value neuron steadily increases as the agent approaches reward. Moreover, Frémaux et al. (2013) discuss an interesting link between their simulations and the behavior of real "ramp" neurons studied by van der Meer and Redish (2011). Therefore, we believe that the ramp intuition deserves further analysis, to better understand its potential use as a learning principle.

Our aim in this paper is to study the value-ramp principle from a general reinforcement learning perspective. Thereto, we formalize the intuition with a concrete algorithm, called Value-Ramp. As in Q-learning (Watkins, 1989; Watkins and Dayan, 1992), we compute a value $V(s,a)$ for each state-action pair (s,a). The state value $V[s]$ is the maximum across the actions, i.e., $V[s] = \max\{V(s,a) \mid a \text{ is an action}\}$. Letting $R(s,a)$ be a nonnegative reward quantity obtained when performing a in s, and letting s' be the successor state, Value-Ramp updates $V(s,a)$ as follows:

$$V'(s,a) := \max(0, V(s,a) + d), \tag{1.1}$$

where

$$d = \max(V[s'], R(s,a)) - K - V[s],$$

and $K \geq 1$ is a fixed step size. Rewards are assumed to be nonnegative, and values are constrained to be nonnegative.[1] We assume throughout this paper that values are natural numbers; natural numbers are adequate for our study. As a benefit, natural numbers can be implemented compactly and robustly on a computer.

Intuitively, value is reward expectation, or closeness to reward. On a path, Value-Ramp propagates encountered reward quantities (and values) backwards in time, where each step subtracts K. The effect is that when following the path forwards, we see increasingly larger values, and there is a reward peak at the end. By choosing actions to maximize value, the agent can follow an increasing ramp of values. After learning, the agent experience may appear as in Figure 1.1.

[1] The nonnegative range is inspired by biological learning models where the (positive) reward spectrum has a dedicated representation mechanism, leaving room for a dual mechanism to represent the aversive spectrum (Schultz, 2013; Hennigan et al., 2015). Representing the aversive spectrum is left as an item for further work (see Section 5).

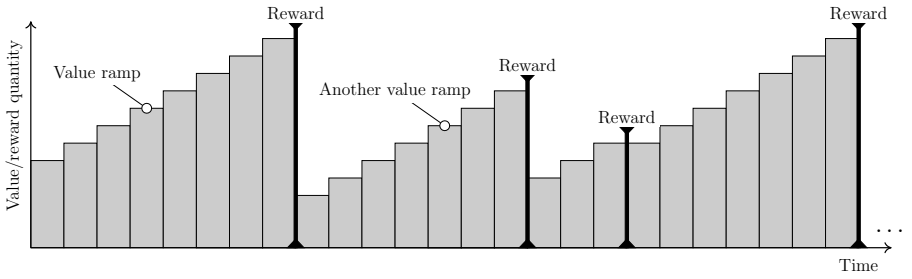

Figure 1.1: Illustration of the ramp-like value experience of the agent.

To find concrete insights about the resulting agent behavior, our approach is to formally study Value-Ramp on well-defined tasks. This approach can be likened to devising specific experiments in which the observed agent behavior is described. An important difference, though, is that we formally prove the observations. Although one may expect real tasks to be more complex than those investigated here, it still appears beneficial to have concrete insights about well-defined circumstances. Possibly, real-world behaviors can be understood as a mixture of formally described behaviors. Below, we summarize the insights of our study in an informal manner.

Exploration (Theorem 3.1 and Theorem 3.10)**.** When exploring, the agent wields a global viewpoint where it can reason about different reward magnitudes in the task. The agent could compute a height-map of values. To elaborate that intuition, we have considered deterministic tasks, where the agent always sees the same successor state for each state-action pair. We additionally assume that any state can be reached from any other state. For each state s, reward quantities become less important when they are remote from s. We show the following: by repeatedly trying all state-action pairs (in exploratory fashion), the agent learns for each state which rewards have the best quantity-versus-distance trade-off. Subsequently, by choosing actions to maximize value, the agent continuously moves to the highest reward as soon as possible. This insight shows the potential use of Value-Ramp as a behavior optimizer.

Greediness (Theorem 4.6)**.** When constantly choosing actions to maximize reward, i.e., in a greedy approach, the agent has a local viewpoint restricted to measuring progress along a path. Here, the agent does not care about all rewards, just about finding one reward. This is useful for navigational tasks, even in abstract state spaces. To elaborate that intuition, we have considered nondeterministic tasks, where the successor state resulting from a state-action pair could vary. We make the relaxing assumption that the state space can be viewed as a stack of layers, where the bottom layer contains reward, and where the states at each layer can move robustly into a deeper layer (but without knowing the precise successor). We show the following: by constantly choosing actions to maximize value, the agent eventually learns to completely avoid cycles without reward. Phrased differently, eventually, whenever the agent walks in a cycle, the cycle is broken by reward (no matter how small the cycle is). This insight shows that Value-Ramp can keep navigating to reward, even in tasks with a degree of unpredictability.

The above insights apply to many tasks, ranging from navigation on maps to finding rewarding strategies in abstract state spaces. Interestingly, Value-Ramp appears easy to configure. The theorems work for any step size $K \geq 1$, but in practice one could simply take $K = 1$. We also introduce a parameter ϵ to control the degree of exploration, which is common practice in reinforcement learning. Other approaches in reinforcement learning often have multiple parameters, e.g., in Q-learning (Watkins, 1989; Watkins and Dayan, 1992) one has the learning rate and the reward discounting factor.

In summary, the Value-Ramp algorithm has useful characteristics: (1) it is conceptually simple, (2) it is easy to configure, and (3) it has a stable implementation based on natural numbers. Additionally, insights discussed in this paper suggest that the algorithm might be versatile.

Outline This paper is organized as follows. Section 2 formalizes the Value-Ramp algorithm and tasks. Next, Section 3 contains the insights about exploration and optimization on deterministic tasks. Section 4 contains the insight about greedy learning on nondeterministic tasks. We conclude with items for further work in Section 5.

2 Value-Ramp algorithm

In this section, we introduce the Value-Ramp algorithm in a general reinforcement learning setting (Sutton and Barto, 1998). In subsequent sections, we analyze the behavior of the algorithm on two classes of applications, one based on continued exploration (Section 3) and the other based on greedy path following (Section 4).

2.1 Basic definitions

Suppose we have a finite set S of states and a finite set A of actions. We write

$$s \xrightarrow{a} s'$$

to denote that we can reach state s' by applying action a to state s.

As in Q-learning (Watkins, 1989; Watkins and Dayan, 1992), we assign a numerical value to each pair $(s, a) \in S \times A$. This setup reflects the intuition that states by themselves do not necessarily have meaning, but rather it is the intention, or action, in the state that matters. In the present paper, natural numbers are sufficient for representing values. Hence, a value function V is of the form

$$V : S \times A \to \mathbb{N}.$$

We define the value of a state s, denoted $V[s]$, as the maximum of the values over the actions:

$$V[s] = \max \{V(s, a) \mid a \in A\} .$$

The set of actions preferred by s in V, denoted $\mathit{pref}(s, V)$, contains the actions with the highest value in s:

$$\mathit{pref}(s, V) = \{a \in A \mid V(s, a) = V[s]\} .$$

Note that always $pref(s, V) \neq \emptyset$.

For each pair $(s, a) \in S \times A$, we have an immediate reward quantity $R(s, a) \in \mathbb{N}$ to say how good action a is in state s. The reward is given externally to the agent, whereas a value function forms the internal belief system of the agent about expectations (of reward).

As convenience notation, for each integer $x \in \mathbb{Z}$, we define a clamping operation

$$[\![x]\!] = \begin{cases} x & \text{if } x \geq 0; \\ 0 & \text{if } x < 0. \end{cases}$$

For any two integers x and y, note that $x \leq y$ implies $[\![x]\!] \leq [\![y]\!]$. In the proofs we also frequently use the equality $\max([\![x]\!], [\![y]\!]) = [\![\max(x, y)]\!]$.[2]

2.2 Desired properties

Suppose we have a path

$$s_1 \xrightarrow{a_1} \ldots \xrightarrow{a_{n-1}} s_n \xrightarrow{a_n} s_{n+1}.$$

We fix some step size $K \in \mathbb{N}$ with $K \geq 1$. Now, if the agent would repeatedly visit the above path, our intention of the Value-Ramp algorithm is to find a value function V with the following properties: for each $i \in \{1, \ldots, n\}$,

- for each $j \in \{i, \ldots, n\}$, we have $V(s_i, a_i) \geq [\![R(s_j, a_j) - (j - i + 1)K]\!]$; and,
- $V(s_i, a_i) \leq [\![\max(V[s_{i+1}], R(s_i, a_i)) - K]\!]$.

The first property says that value should reflect reward expectation, taking into account the time until reward, as illustrated in Figure 2.1.[3] Rewards can be surpassed by value expectations for larger rewards, as illustrated in Figure 2.2. The first property suggests a mechanism for increasing value, to propagate reward information backwards in time.

The second property says that local maxima can be sustained only by reward; other local maxima should be gradually removed, as illustrated in Figure 2.3. We might refer to local value maxima without reward as violating values. The second property suggests a mechanism for subtracting value, in order to maintain the ramp shape.

Remark 2.1 (Note on violating values)**.** In absence of reward, the only way to prevent nonzero values from being labeled as violating, is to have an infinite ramp of increasingly larger values. But there are only a finite number of states, so eventually the ramp should meet true reward. □

Remark 2.2 (Outlook)**.** A value ramp reflects some ideal value function that, likely, can only be obtained under the right circumstances. Moreover, it might be difficult to describe properties that are both interesting and sufficiently general, because there are widely different kinds of tasks (or environments) upon which

[2] If $x \leq y$ then $[\![x]\!] \leq [\![y]\!]$; subsequently, $\max([\![x]\!], [\![y]\!]) = [\![y]\!] = [\![\max(x, y)]\!]$. The other case is symmetrical.

[3] If $j = i$ then we subtract K once from $R(s_i, a_i)$, to indicate that action a_i should first be executed at state s_i before the reward is given.

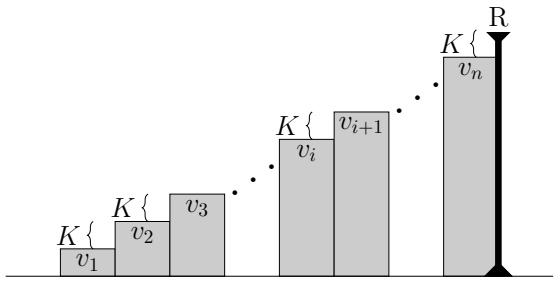

Figure 2.1: Depiction of a value ramp. The agent is following a path of n states, whose values $v_1, \ldots, v_n$ (in the encountered value functions) form a ramp. For graphical simplicity, there is only one nonzero reward, at the end. The height of the reward "pillar" reflects the actual reward quantity.

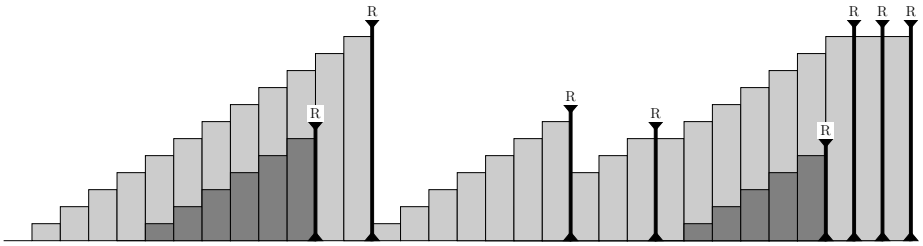

Figure 2.2: Rewards can be surpassed by higher value expectations. We have given a darker shade to the ramps of surpassed rewards. Also, note that in principle the same reward quantity could be repeated in subsequent time steps (suggested at the end of this figure).

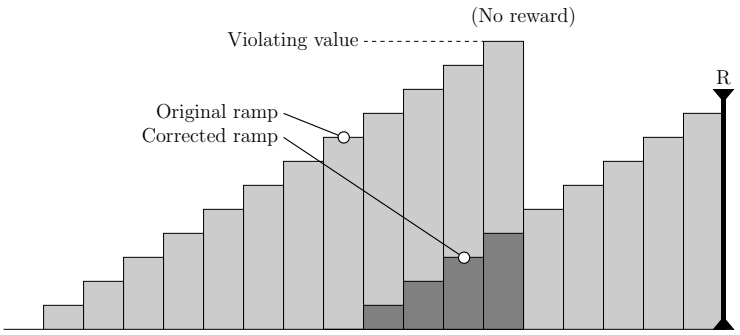

Figure 2.3: To obtain ramps in alignment with reward, we should remove local maxima without reward, called violations.

an agent could operate. For these reasons, in Section 3 and Section 4 we provide more detailed insights for specific classes of tasks. This could provide an initial foundation for understanding the value ramp principle. □

2.3 Concrete algorithm

A learning rule is a function that produces a value change (as an integer) when given a triple $(v, v', r) \in \mathbb{N} \times \mathbb{N} \times \mathbb{N}$, where v is the value of the current state, v' is the value of the next state, and r is the reward quantity observed during the transition from the current state to the next state. The desired properties from Section 2.2 inspire a concrete learning rule Δ_K, where $K \geq 1$, defined for each triple (v, v', r) as

$$\Delta_K(v, v', r) = \max(v', r) - K - v. \tag{2.1}$$

The proposed value change could be either strictly positive, zero, or strictly negative.

Remark 2.3 (Usage of state value)**.** Recall that state value is defined as the maximum over state-action values. There are multiple reasons for using state values to compute the update in Equation (2.1) instead of using bare state-action values. First, using state value appears more biologically plausible, since the brain likely assigns one (global) value to each observed state (Potjans et al., 2011; Frémaux et al., 2013). When moving from state to state, the global value could be the aggregate of detailed state-action values. Second, if we try to use bare state-action values instead, there is no natural mechanism to select which action value to associate with an observed successor state; causing us to default to the highest action value for instance. □

Based on Equation (2.1), we now formalize how the agent updates value through experience. A configuration of the system is a pair (s, V) saying that we are in state s and that the current value function is V.

Definition 2.1 (Transition)**.** A transition is a quintuple $(s_1, V_1, a_1, s_2, V_2)$, where (s_1, V_1) and (s_2, V_2) are configurations; s_2 is reached from s_1 through action a_1; and, V_2 is defined, for each $(s, a) \in S \times A$, as

$$V_2(s,a) = \begin{cases} [\![V_1(s_1, a_1) + \Delta_K\big(V_1[s_1], V_1[s_2], R(s_1, a_1)\big)]\!] & \text{if } (s,a) = (s_1, a_1); \\ V_1(s,a) & \text{otherwise.} \end{cases}$$

We emphasize that during the transition, the value of the successor state s_2 is based on the old value function V_1. We also write the transition as

$$(s_1, V_1) \xrightarrow{a_1,\, s_2} (s_2, V_2).$$

Algorithm 1 gives pseudocode for performing transitions. The full Value-Ramp algorithm, shown in Algorithm 2, repeatedly generates transitions. Note that there is a probability ϵ at each time step of choosing from all actions, instead of choosing from the actions with highest value. If $\epsilon = 0$ then the algorithm follows the best known path to reward, without further exploration; in that case, we say that the algorithm is greedy.

Algorithm 1 Value-Ramp update function.
Input:
- (s, V): current configuration
- a: performed action
- s': observed successor state

function UPDATE(s, V, a, s')
  $d := \max(V[s'], R(s, a)) - K - V[s]$
  Define V' as V but set $V'(s, a) := [\![V(s, a) + d]\!]$
  return V'
end function

Algorithm 2 Value-Ramp algorithm.
Input:
- V_1: initial value function (random)
- s_1: initial start state
- K: step size with $K \geq 1$
- ϵ: probability in $[0, 1]$

procedure VALUE-RAMP
  $V := V_1$
  $s := s_1$
  repeat
    $a :=$ choose from $pref(s, V)$
    With probability ϵ, do $a :=$ choose from A
    $s' :=$ some state resulting from (s, a)
    $V' :=$ Update(s, V, a, s') ▷ See Algorithm 1
    $s := s'$
    $V := V'$
  until Interrupt
end procedure

Remark 2.4 (Natural numbers)**.** Since we use a discrete time framework, natural numbers are a perfect fit for representing the steps of a ramp. Practical implementations of natural numbers are robust under addition and subtraction. Also, as is commonly known, a string of n bits can represent any natural number in the range $\{0, \ldots, 2^n - 1\}$. Modest storage requirements can therefore accommodate huge values. That might be useful for learning (very) long paths in navigation problems (see Section 4).

Many approaches in reinforcement learning are based on rational numbers (Sutton and Barto, 1998). Approximation errors arise when rational numbers are implemented as floating point numbers, inspiring the development of new digital number formats (Gustafson, 2015). By using natural numbers, Value-Ramp avoids approximation errors. □

Remark 2.5 (Parameters)**.** A first parameter of Value-Ramp is the step size K. We develop the formal insights for a general $K \geq 1$ (e.g. Theorem 3.1 and Theorem 4.6). In practice, it might be useful to simply set $K = 1$, because then rewards generate longer ramps, allowing the agent to learn longer strategies

to reward. Second, the exploration probability ϵ in Algorithm 2 is a standard principle in reinforcement learning (Sutton and Barto, 1998).

Value-Ramp has no other parameters besides K and ϵ. In comparison, the general framework of reinforcement learning introduces an α and γ parameter (Sutton and Barto, 1998). This applies in particular to Q-learning (Watkins, 1989; Watkins and Dayan, 1992), which has famous applications (Mnih et al., 2015). Parameter α can be understood as the learning rate. Parameter γ, representing reward-discounting, is slightly less intuitive and could require detailed knowledge of the task domain in order to produce desired agent behavior (Schwartz, 1993).

Value-Ramp replaces the α parameter by a fast value update mechanism that immediately establishes a (local) ramp shape on encountered states. An item for further work is to slow down the value update in the context of biological plausibility (see Section 5).

Value-Ramp dismisses the γ parameter by directly using reward quantities to define the height of ramps. For each reward, the ramp shape establishes a natural trade-off between the quantity of a reward and the time to get there. An item for further work is to investigate in more detail the relationship between reward discounting and the value ramp principle (see Section 5). □

Remark 2.6 (Fixing K)**.** All definitions and results hold for any $K \geq 1$. But for notational simplicity, we choose not to mention the symbol "K" explicitly in the notations. We assume that throughout the rest of the paper, some particular $K \geq 1$ is fixed. □

2.4 Tasks formalized

We want to describe the effect of Value-Ramp on tasks. Formally, a task T is a tuple $(S, S^{\text{start}}, A, \delta, R)$, where

- S is a finite, nonempty, set of states;
- $S^{\text{start}} \subseteq S$ is a set of start states;
- A is a finite, nonempty, set of actions;
- $\delta : S \times A \to \mathcal{P}(S)$ is the transition function, where $\delta(s, a) \neq \emptyset$ for each $(s, a) \in S \times A$;[4] and,
- $R : S \times A \to \mathbb{N}$ is the reward function.

For any $(s, a, s') \in S \times A \times S$, we write $s \xrightarrow{a} s'$ if $s' \in \delta(s, a)$.

A run of Value-Ramp on T is an infinite sequence of transitions, where the target configuration of each transition is the source configuration of the next transition,

$$(s_1, V_1) \xrightarrow{a_1,\, s_2} (s_2, V_2) \xrightarrow{a_2,\, s_3} (s_3, V_3) \xrightarrow{a_3,\, s_4} \ldots,$$

where $s_1 \in S^{\text{start}}$, and for each $i \geq 1$ we have $s_{i+1} \in \delta(s_i, a_i)$. For each $i \geq 1$, we recall that V_{i+1} is uniquely determined by (s_i, V_i, a_i, s_{i+1}) (see Section 2.3). We allow V_1 to be a random value function. We emphasize that the successor state of each transition is restricted by function δ.

[4]For a set X, the symbol $\mathcal{P}(X)$ denotes the powerset of X, which is the set of all subsets of X.

The following lemma is a general observation that we will use frequently in proofs:

Lemma 2.7. For any task, in any infinite transition sequence, there are only a finite number of possible configurations.

Proof. Let $(S, S^{\text{start}}, A, \delta, R)$ be the task. For a value function V we define $\mathit{ceiling}(V) = \max(m_1, m_2)$ where

$$m_1 = \max\{R(s,a) \mid (s,a) \in S \times A\}\text{; and,}$$
$$m_2 = \max\{V(s,a) \mid (s,a) \in S \times A\}.$$

Intuitively, $\mathit{ceiling}(V)$ is the highest quantity accessible by the agent; this quantity is either defined by reward or by the value function itself. For each transition

$$(s, V) \xrightarrow{a, s'} (s', V'),$$

we can show that $\mathit{ceiling}(V) \geq \mathit{ceiling}(V')$ (see Appendix A). By transitivity, for every infinite transition sequence, the ceiling quantity of the first value function is an upper bound on the ceiling quantity of all subsequent value functions. So, the infinite transition sequence has a finite number of value functions because (1) there is an upper bound on the values, (2) value functions are composed of natural numbers, and (3) there are a finite number of states and actions. Therefore there are a finite number of configurations. □

Remark 2.8 (Perception and finiteness)**.** The task structure represents how the agent perceives its environment. The agent perception is in general the result of various processing steps applied to sensory information. Agent perception is not the focus of this paper. Although the environment in which the agent resides could have infinitely many states, we assume that the agent has a limited conceptual framework consisting of finitely many states. We still allow many states though. The finiteness of the state space is important for the convergence proofs of this paper; more precisely, the assumption is used in the general Lemma 2.7. □

2.4.1 Kinds of run: exploring versus greedy

Hereafter, we restrict attention to two kinds of run.

Exploring First, we say that a run is exploring if the following holds: if a configuration (s, V) occurs infinitely often in the run, then for each $a \in A$ and each $s' \in \delta(s, a)$, there are infinitely many transitions

$$(s_i, V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1}, V_{i+1}),$$

where $(s_i, V_i, a_i, s_{i+1}) = (s, V, a, s')$. Intuitively, an exploring run contains a fairness assumption to ensure that the system explores infinitely often those options that are infinitely often available.

Greedy Second, we say that a run is greedy if the following holds:

1. each transition $(s, V) \xrightarrow{a,\, s'} (s', V')$ in the run satisfies $a \in \mathit{pref}(s, V)$; and,
2. if a configuration (s, V) occurs infinitely often in the run, then for each $a \in \mathit{pref}(s, V)$ and each $s' \in \delta(s, a)$, there are infinitely many transitions

$$(s_i, V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1}, V_{i+1}),$$

where $(s_i, V_i, a_i, s_{i+1}) = (s, V, a, s')$.

In a greedy run, we always select a preferred action, but the system can not reliably choose only one action from equally-preferred actions; moreover, as a fairness assumption, the system can not indefinitely postpone witnessing a certain successor state.

Remark 2.9 (Relationship with Algorithm 2)**.** In Algorithm 2, we generate exploring runs by setting $\epsilon > 0$. We will not use the specific ϵ value to delineate strict subclasses of exploring runs whose exploration rate satisfies ϵ. In Algorithm 2, we generate greedy runs by setting $\epsilon = 0$.

When running Algorithm 2 on a task $T = (S, S^{\text{start}}, A, \delta, R)$, we assume that if the same state-action pair (s, a) is executed infinitely often then each successor state in $\delta(s, a)$ is infinitely often the result of (s, a). □

3 Exploration on deterministic tasks

In a first study, we would like to show optimal value estimation of Value-Ramp on at least some (well-behaved) class of tasks. Thereto we consider tasks that are both deterministic and connected, abbreviated DC. In Section 3.1 we show that exploring runs learn optimal values on DC tasks. In Section 3.2, we subsequently show that when the agent uses the optimal values to select actions, the agent follows so-called optimal paths. In Section 3.2.1, we apply the results to shortest path following.

3.1 Optimal value estimation

We first define a few auxiliary concepts. Let $T = (S, S^{\text{start}}, A, \delta, R)$ be a task. To improve readability, we omit symbol T from the notations below where possible; it will be clear from the context which task is meant.

DC tasks We say that T is deterministic if $|\delta(s, a)| = 1$ for each $(s, a) \in S \times A$. Next, we say that T is connected if for each $(s, s') \in S \times S$, there is a path

$$s_1 \xrightarrow{a_1} \ldots \xrightarrow{a_{n-1}} s_n$$

with $s_1 = s$ and $s_n = s'$. Connectedness means that for each state we can go to any other state. We say that a task is DC if the task is both deterministic and connected.

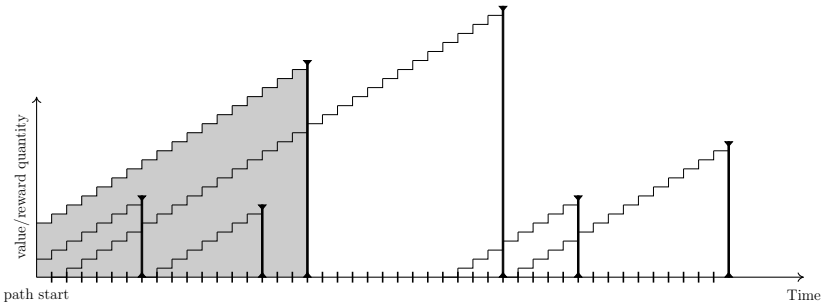

Figure 3.1: Illustration of action-path value. Each tick mark on the horizontal axis represents a state-action pair of the action-path. To keep the figure simple, only some of the state-action pairs have a strictly positive reward, leading to a value-ramp. We have shaded the value-ramp of the state-action pair that determines the path value. The other state-action pairs project less reward expectation towards the beginning of the path.

Consistency On deterministic tasks, it will be interesting to observe eventual stability of the value function. In that context, we say that a value function V is consistent if it satisfies: $\forall s \in S$, $\forall a \in \mathit{pref}(s, V)$, denoting $\delta(s, a) = \{s'\}$,

$$V[s] = [\![\max(V[s'], R(s, a)) - K]\!] .$$

Intuitively, this means that the agent knows exactly what value to expect when following preferred actions. We will see below in the context of Corollary 3.2 that consistency eventually halts the learning process on DC tasks.

Optimal value Next, we define a notion related to shortest paths. Let s be a state. An action-path for s is a sequence

$$p = (s_1, a_1), \ldots, (s_{n-1}, a_{n-1}), (s_n, a_n),$$

of state-action pairs, where $s_1 = s$, and $s_{i+1} \in \delta(s_i, a_i)$ for each $i \in \{1, \ldots, n-1\}$. We allow $n = 1$. We define the value of p, denoted $\mathit{val}(p)$, as

$$\mathit{val}(p) = \max \{[\![R(s_i, a_i) - iK]\!] \mid i \in \{1, \ldots, n\}\} . \tag{3.1}$$

This value expresses a trade-off between time and reward amplitude. For example, high rewards could become less important than lower rewards if the time distance is too long. The concept is illustrated in Figure 3.1. Note that always $\mathit{val}(p) \geq 0$ due to the clamping operation.

If the action-path p contains a cycle of states then there is always an action-path p' without such cycles and with $\mathit{val}(p') \geq \mathit{val}(p)$. To see this, we can do the following steps to transform p into a cycle-free action-path without decreasing the value:

1. We select some $i \in \{1, \ldots, n\}$ with $\mathit{val}(p) = [\![R(s_i, a_i) - iK]\!]$.
2. We remove all pairs (s_j, a_j) with $j > i$.

3. In the remaining path, we systematically replace all cycles $(s, a), \ldots, (s, a')$ (with repeated state s) by the single step (s, a'). Note that pair (s_i, a_i) is preserved because this pair comes last, as caused by step 2. As a result, the reward quantity $R(s_i, a_i)$ can only come closer to the beginning of the path.

Let *explore*(s) be the set of all cycle-free action-paths starting at state s. We define the optimal value of s, denoted *opt-val*(s), as

$$\textit{opt-val}(s) = \max \{ \textit{val}(p) \mid p \in \textit{explore}(s) \},$$

i.e., the optimal value is the largest value across the (cycle-free) action-paths. The case *opt-val*$(s) = 0$ occurs when all reward is too remote for s.

We say that a value function V is optimal if it satisfies: $\forall s \in S$,

$$V[s] = \textit{opt-val}(s).$$

We are now ready to state the optimization result:

Theorem 3.1 (Optimization)**.** For each DC task, in each exploring run, eventually every value function is both optimal and consistent.

The proof is given in Section 3.3. The following corollary provides an additional insight about the learning process on DC tasks:

Corollary 3.2. For each DC task, in each exploring run, eventually the value function is no longer changed, i.e., there is a fixpoint on the value function.

The proof is given in Section 3.4.

Example 3.3 (Example simulation)**.** To illustrate Theorem 3.1, we have simulated the Value-Ramp algorithm on a 2D grid world that is both deterministic and connected (DC). Each cell $(x, y) \in \mathbb{N} \times \mathbb{N}$ inside the boundaries of the map is a distinct state. There is a fixed start cell. At each cell, there are five deterministic actions available to the agent: left, right, up, down, and finish. The agent can not move through wall cells, serving as obstacles. Some cells are marked as goal cells. By performing the finish action in a goal cell g, the agent receives a fixed reward quantity associated with goal cell g, and the agent is subsequently sent back to the fixed start cell. In a non-goal cell, the finish action neither gives reward and neither moves the agent to another cell. The agent learns the values of all cell-action pairs. Figure 3.2 shows for three different maps how value is propagated from the goal cells across the map. Eventually, the cell values visibly stabilize; one could imagine that this is the point after which the value function (1) is optimal and consistent (Theorem 3.1) and (2) no longer changes (Corollary 3.2). □

Remark 3.4 (Degree of exploration)**.** When relating exploring runs to Algorithm 2, we would like to point out that Theorem 3.1 works for any $\epsilon > 0$, even for very small (but nonzero) ϵ values that would make the agent seem almost entirely greedy. Therefore, the theorem might be useful for better understanding settings where a high degree of greediness (and therefore exploitation of knowledge) is preferred. □

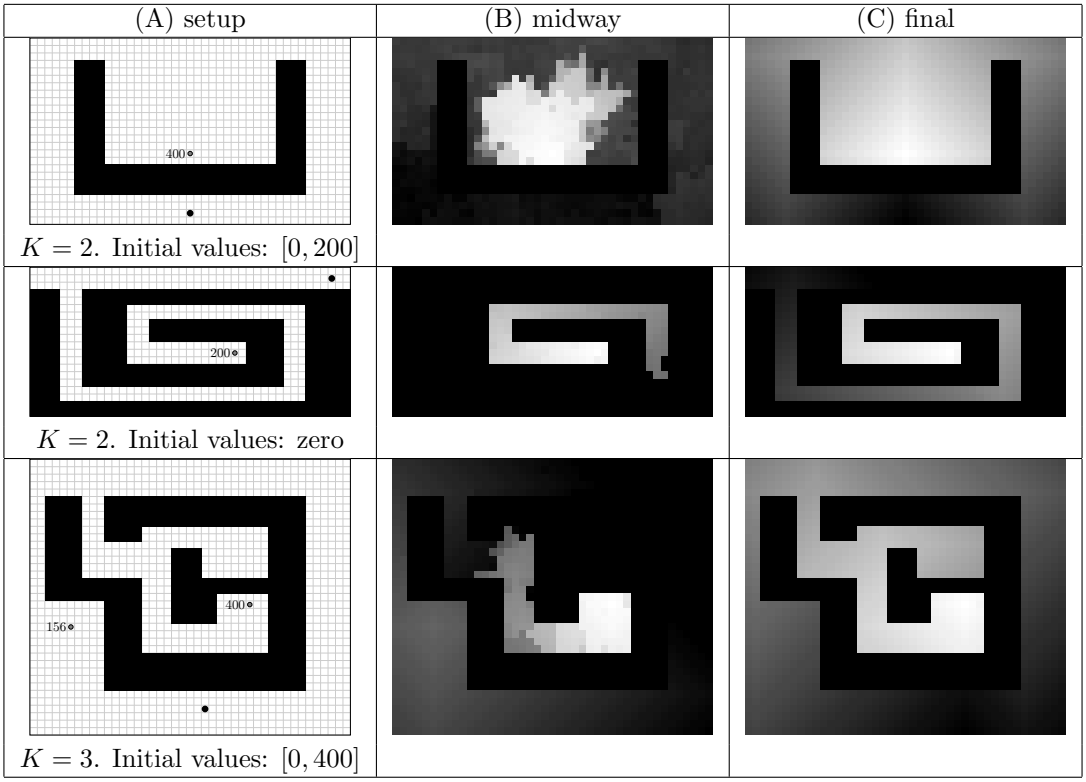

Figure 3.2: Simulation result of Value-Ramp on three 2D grid maps, that are deterministic and connected (see description in Example 3.3). To have faster convergence, we considered $\epsilon = 1$.
(A) Column A shows the setup of each map. The first and second map use random initial values (for each state-action pair) in the mentioned interval. The starting location is indicated by a black dot, and the goal cells are marked with their numerical reward quantity. We also vary K to test more circumstances.
(B) For a value function V midway the learning process, column B shows the value $V[c]$ of each cell c, computed as the maximum over the action values for that cell. The highest values are shown as the brightest.
(C) Column C shows the cell value when no visual changes occur anymore. Note that in the second row (with the spiral), some top-right cells converge to zero (optimal) value.

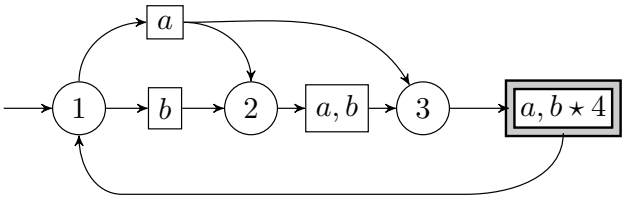

Figure 3.3: A task, where states are represented by circles, and action applications are represented by rectangles. Start states have an inbound arrow without origin. The nonzero rewards are indicated by a shaded box. In this case, $(3, a)$ and $(3, b)$ are assigned a reward quantity of 4.

Remark 3.5 (Liberal initialization)**.** Theorem 3.1 applies to exploring runs that start with arbitrarily initialized value functions. This highlights a strength of the Value-Ramp algorithm. In particular, the theorem seems to refute immediate simplifications of Value-Ramp that would simply remember for each state-action pair the highest value seen so far. Such a simplification would in general require that initial values are all zero, which is not needed by Theorem 3.1. □

Remark 3.6 (Off-policy learning)**.** Theorem 3.1 resembles the viewpoint of Q-learning (Watkins, 1989; Watkins and Dayan, 1992) in that the agent is updating its value estimation without necessarily using its learned knowledge to explore. Essentially, all that we require in the proof is that the system keeps running, and that each state-action pair is visited sufficiently often. This has been called off-policy learning by Sutton and Barto (1998): the agent is trying to find an optimal policy (mapping states to the best actions), independently of the other policy used to explore the state space. □

Remark 3.7 (Not all equivalent actions)**.** In a DC task, in an exploring run, we can expect that the agent only rarely learns two optimal actions for the same state. Once the agent has found one optimal action a for a state s, it will become more difficult (or impossible) to increase the value for another pair (s, a'). Indeed, if state s has reached its optimal value, through the value of (s, a), the value of s has become too high to have positive surprise when trying the pair (s, a'); positive surprise would correspond to $d > 0$ in Algorithm 1. □

On nondeterministic tasks, the following example illustrates why exploring runs do not necessarily converge numerically (as in Theorem 3.1).

Example 3.8 (Nondeterminism causes fluctuations)**.** We consider the nondeterministic task in Figure 3.3. Note that the state-action pair $(1, a)$ can choose among two successors: 2 and 3. For states 2 and 3, the actions behave deterministically. For simplicity, we assume that $K = 1$ and that initial values are zero. In an exploring run, starting at state 1, the pairs $(3, a)$ and $(3, b)$ both get the value $4 - K = 3$; subsequently, the pairs $(2, a)$ and $(2, b)$ get the value $3 - K = 2$. Since $(1, b)$ always arrives at state 2, the value of $(1, b)$ will also stabilize at 1. However, $(1, a)$ will continue to fluctuate in value: if $(1, a)$ arrives at state 2 then the value will be 1, and if $(1, a)$ arrives at state 3 then the value will be 2. □

Later, in Section 4, we will approach nondeterministic tasks with greedy runs instead of exploring runs; and, we will restrict attention to so-called navigation problems, that are introduced in Section 3.2.1.

3.2 Optimal path following

The previous Section 3.1 was about learning optimal values. Here, we study the effect of value optimization on the actual behavior of the agent.

Definition 3.1. For a given task, we say that a run fragment

$$(s_1, V_1) \xrightarrow{a_1,\, s_2} \dots \xrightarrow{a_{n-1},\, s_n} (s_n, V_n) \xrightarrow{a_n,\, s_{n+1}} (s_{n+1}, V_{n+1}),$$

where $n \geq 1$, is a value-sprint if

1. $V_i[s_i] \leq V_i[s_{i+1}] - K$ for each $i \in \{1, \dots, n-1\}$; and,
2. $V_n[s_n] > V_n[s_{n+1}] - K$.

In a value-sprint, each transition witnesses increasingly larger values, separated by at least K, except the last transition. A value-sprint could occur anywhere in the run (not necessarily at the beginning). We allow $n = 1$, in which case condition (1) is vacuously true. Note that we can not split a value-sprint in smaller value-sprints: the first part would not be a value-sprint because condition (2) is not satisfied.

The following lemma relates runs and value-sprints:

Lemma 3.9. For each task, every run is always an infinite sequence of value-sprints.

Proof. Let $T = (S, S^{\text{start}}, A, \delta, R)$ be the task. Suppose towards a contradiction that there is a run $\mathcal{X}$ that is not an infinite sequence of value-sprints. Then $\mathcal{X}$ is a finite sequence of value-sprints followed by an infinite tail

$$(s_i, V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1}, V_{i+1}) \xrightarrow{a_{i+1},\, s_{i+2}} \dots,$$

where $V_j[s_j] \leq V_j[s_{j+1}] - K$ for each $j \geq i$.

Let $j \geq i$. We note that $V_j[s] \leq V_{j+1}[s]$ for each $s \in S$: in Algorithm 1, we can use the assumption $V_j[s_j] \leq V_j[s_{j+1}] - K$ to obtain

$$\begin{aligned} d &= \max(V_j[s_{j+1}], R(s_j, a_j)) - K - V_j[s_j] \\ &\geq V_j[s_{j+1}] - K - V_j[s_j] \\ &\geq 0. \end{aligned}$$

We observe that $V_j[s_j] < V_{j+1}[s_{j+1}]$: everything combined, we have $V_j[s_j] < V_j[s_j] + K \leq V_j[s_{j+1}] \leq V_{j+1}[s_{j+1}]$. By transitivity, for any indices j and k with $i \leq j < k$, we have $V_j[s_j] < V_k[s_k]$. But then we would encounter infinitely many (state) values, and thus infinitely many configurations, contradicting Lemma 2.7. □

The above value-sprint describes the following action-path, where we omit the last state s_{n+1}:

$$p = (s_1, a_1), \dots, (s_n, a_n).$$

We say that p is optimal (for s_1) if $\mathit{val}(p) = \mathit{opt\text{-}val}(s_1)$. We are now ready to express the effect of value optimization on the behavior of the agent:

Theorem 3.10 (Follow optimal paths)**.** For each DC task, for each greedy run that starts with an optimal and consistent value function, each value-sprint describes an optimal action-path.

The proof is given in Section 3.5.

Remark 3.11 (Increasingly better)**.** At moments when the agent is not exploring, and is greedily applying preferred actions, the agent is following its best guess about optimal paths. Although we do not know the precise moment when the agent has complete knowledge about optimal paths, we can imagine that the agent is increasingly getting better at following them. Theorem 3.1 tells us that the value function will eventually contain the knowledge about optimal paths. Then, by Theorem 3.10, any subsequent greedy fragments of the run follow optimal paths. Note that parameter ϵ determines the amount of time that the agent exploits its knowledge; high values for ϵ could make the agent still seem random, even if the agent has knowledge of optimal paths. □

3.2.1 Shortest paths

As an application and further explanation of Theorem 3.1 and Theorem 3.10, we discuss a relationship between path value and shortest paths. See (Cormen et al., 2009) for an introduction to the shortest path problem and related algorithms. The standard shortest-path algorithms process graph data in bulk fashion, e.g., they can iterate over vertices and edges. A reinforcement learning system, on the other hand, builds its belief by (repeatedly) following trajectories through the transition function of the task.

We first consider the following definition.

Definition 3.2. We say that a task $T = (S, S^{\text{start}}, A, \delta, R)$ is a navigation problem if there is exactly one nonzero reward quantity M, and with $M > |S|\, K$. More precisely, for all $(s, a) \in S \times A$ we have either $R(s, a) = M$ or $R(s, a) = 0$, and there is at least one (s, a) with $R(s, a) = M$.

In a navigation problem, the intention behind the sufficiently large reward quantity is to allow the agent to learn a (cycle-free) path between any two states.

DC navigation problems A DC navigation problem is a navigation problem that is also deterministic and connected. Let $T = (S, S^{\text{start}}, A, \delta, R)$ be a DC navigation problem. Let $p = (s_1, a_1), \ldots, (s_n, a_n)$ be an action-path. We say that p is rewarding if there is at least one $i \in \{1, \ldots, n\}$ with $R(s_i, a_i) = M$. If p is rewarding then we define the length of p, denoted $\mathit{len}(p)$, as the smallest $i \in \{1, \ldots, n\}$ with $R(s_i, a_i) = M$. Note that if a rewarding action-path contains cycles then we can transform it into a rewarding action-path without cycles, using the procedure at the beginning of Section 3.1.

Thanks to connectedness, there is a cycle-free action-path from each state s to a state-action pair (s', a') with $R(s', a') = M$. Recalling the definition of path value from Equation (3.1), cycle-free rewarding action-paths have a strictly positive value because $M > |S|\, K$: on a cycle-free path, each pair (s_i, a_i) contributes at least $[\![R(s_i, a_i) - |S|\, K]\!]$ to the overall path value; if one of the pairs is rewarding then the overall path value is strictly positive. Hence, each state s has $\mathit{opt\text{-}val}(s) > 0$. We therefore consider DC navigation problems to be solvable from a path finding viewpoint.

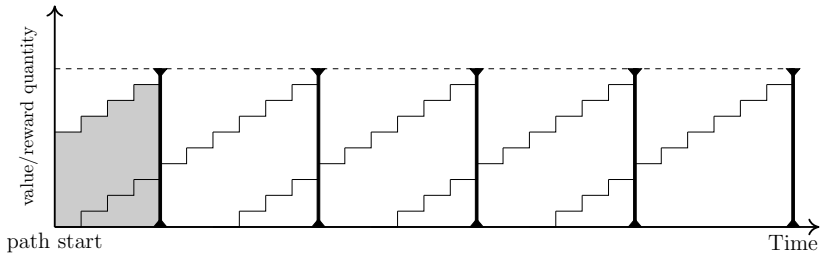

Figure 3.4: When all rewards have the same magnitude, the action-path value is determined only by the distance (or time) to reward. In the figure, we have shaded the ramp of the reward occurrence that determines the path value; it must be the first reward occurrence.

Shortest path following On DC navigation problems, Theorem 3.1 tells us that every exploring run will eventually find an optimal and consistent value function, containing for each state the knowledge of the optimal paths. Theorem 3.10 has the following corollary:

Corollary 3.12 (Follow shortest paths)**.** For all DC navigation problems, for each greedy run that starts with an optimal and consistent value function, each value-sprint follows a shortest path to reward.

Proof. Consider a greedy run starting with an optimal and consistent value function. By Lemma 3.9, the greedy run is an infinite sequence of value-sprints. Moreover, by Theorem 3.10, each value-sprint follows an optimal path. We are left to argue that those paths are actually the shortest (to reward). Since all nonzero reward occurrences have amplitude M, for any two cycle-free rewarding action-paths p_1 and p_2, if $len(p_1) < len(p_2)$ then $val(p_1) > val(p_2)$. This is illustrated in Figure 3.4. Therefore, whenever we follow an optimal action-path p from a state s, we know that p has the shortest length among all those paths that lead from s to reward. □

After the above introduction to navigation problems, we may proceed to the topic of greedy navigation in Section 4.

3.3 Proof of Theorem 3.1

Let us fix some DC task $T = (S, S^{\text{start}}, A, \delta, R)$.

3.3.1 Approach

We first define an auxiliary notion. Let V be a value function. We say that V is valid when for each $(s, a) \in S \times A$, we have

$$V(s, a) \leq [\![\max(V[s'], R(s, a)) - K]\!] .$$

Intuitively, this means that the values are not overestimating the true reward.

We will show in Section 3.3.2 that every exploring run $\mathcal{X}$ has an infinite suffix $\mathcal{X}'$ in which each value function is both valid and optimal. By Property 3.13 (below), all value functions in $\mathcal{X}'$ are also consistent, as desired.

Property 3.13. Let V be a value function. If V is valid and optimal then V is consistent.

Proof. Let $s \in S$ and $a \in \mathit{pref}(s, V)$. Denote $\delta(s, a) = \{s'\}$. We show $V[s] \geq [\![\max(V[s'], R(s, a)) - K]\!]$. Then the validity assumption $V(s, a) \leq [\![\max(V[s'], R(s, a)) - K]\!]$, combined with $V[s] = V(s, a)$ by $a \in \mathit{pref}(s, V)$, implies the desired consistency

$$V[s] = [\![\max(V[s'], R(s, a)) - K]\!] .$$

By Lemma 3.14 (below),

$$\begin{aligned} \mathit{opt\text{-}val}(s) &\geq \max([\![\mathit{opt\text{-}val}(s') - K]\!] , [\![R(s, a) - K]\!]) \\ &= [\![\max(\mathit{opt\text{-}val}(s'), R(s, a)) - K]\!] . \end{aligned}$$

Subsequently, by the optimality assumption on V, which gives $V[s''] = \mathit{opt\text{-}val}(s'')$ for each $s'' \in S$, we have $V[s] \geq [\![\max(V[s'], R(s, a)) - K]\!]$. □

Lemma 3.14. Let $T = (S, S^{\mathrm{start}}, A, \delta, R)$ be a deterministic task. Let $(s, a) \in S \times A$, and denote $\delta(s, a) = \{s'\}$. We always have

$$\mathit{opt\text{-}val}(s) \geq \max([\![\mathit{opt\text{-}val}(s') - K]\!] , [\![R(s, a) - K]\!]).$$

Proof. We have $\mathit{opt\text{-}val}(s) \geq [\![R(s, a) - K]\!]$ because (s, a) is an action-path (of length one) for s. Also, $\mathit{opt\text{-}val}(s) \geq [\![\mathit{opt\text{-}val}(s') - K]\!]$ because any optimal action-path p' for s' can be extended to an action-path for s by adding (s, a) to the front; adding (s, a) to the front pushes the state-action pairs of p' one step further into the future, leading to an overall value decrease with K. □

3.3.2 Obtain validity and optimality

Henceforth, we fix an exploring run $\mathcal{X}$:

$$(s_1, V_1) \xrightarrow{a_1,\, s_2} (s_2, V_2) \xrightarrow{a_2,\, s_3} \dots$$

We show the existence of an infinite suffix $\mathcal{X}'$ in which all value functions are both valid and optimal.

By Property 3.15 (below), we know that in $\mathcal{X}$ we eventually encounter a configuration (s_j, V_j) where V_j is valid and all subsequent value functions are also valid.

Property 3.15. In run $\mathcal{X}$, eventually all encountered value functions are valid. *(Proof in Appendix B.1.)* □

Subsequently, Property 3.16 (below) tells us that after configuration (s_j, V_j), state values do not decrease.

Property 3.16. Consider a transition $(s_i, V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1}, V_{i+1})$. If V_i is valid then for each $s \in S$ we have $V_i[s] \leq V_{i+1}[s]$.

Proof. Let $s \in S$. If $s \neq s_i$ then $V_{i+1}[s] = V_i[s]$. Henceforth, suppose $s = s_i$. If $a_i \notin \mathit{pref}(s_i, V_i)$ then there is some $a' \in \mathit{pref}(s_i, V_i)$ with

$$V_i[s_i] = V_i(s_i, a') = V_{i+1}(s_i, a') \leq V_{i+1}[s_i].$$

Suppose $a_i \in \mathit{pref}(s_i, V_i)$, giving $V_i[s_i] = V_i(s_i, a_i)$. By Algorithm 1,

$$\begin{aligned} V_{i+1}(s_i, a_i) &= [\![V_i(s_i, a_i) + \max(V_i[s_{i+1}], R(s_i, a_i)) - K - V_i[s_i]]\!] \\ &= [\![V_i[s_i] + \max(V_i[s_{i+1}], R(s_i, a_i)) - K - V_i[s_i]]\!] \\ &= [\![\max(V_i[s_{i+1}], R(s_i, a_i)) - K]\!] . \end{aligned}$$

By validity, $[\![\max(V_i[s_{i+1}], R(s_i, a_i)) - K]\!] \geq V_i(s_i, a_i)$. Overall,

$$V_i[s_i] = V_i(s_i, a_i) \leq V_{i+1}(s_i, a_i) \leq V_{i+1}[s_i].$$

□

We summarize what we have so far:

- We eventually reach a configuration (s_j, V_j) where V_j is valid.
- After (s_j, V_j), value functions remain valid and state values do not decrease.

After (s_j, V_j), each state must eventually stop changing its value. Otherwise, since the only change to a state value would be a strict increment, we would see infinitely many state values, and thus infinitely many configurations (contradicting Lemma 2.7). Therefore, somewhere after (s_j, V_j), there is an infinite suffix $\mathcal{X}'$ in which state values no longer change. Let (s_k, V_k) denote the first configuration of $\mathcal{X}'$. In the rest of the proof, we show that V_k is optimal. Hence, all value functions in $\mathcal{X}'$ turn out to be optimal. Overall, all value functions in $\mathcal{X}'$ are both valid and optimal, as desired.

Abbreviate $V = V_k$. Towards a contradiction, suppose V is not optimal. Since V is valid, by Property 3.17 (below) we know that $V[s] \leq \mathit{opt\text{-}val}(s)$ for each $s \in S$. So, if V is not optimal, there is at least one $s \in S$ with $V[s] < \mathit{opt\text{-}val}(s)$.

Property 3.17. Let V be a value function. If V is valid then for each $s \in S$ we have $V[s] \leq \mathit{opt\text{-}val}(s)$. *(Proof in Appendix B.2.)* □

Consider the set

$$\mathit{sub}(V) = \{s \in S \mid V[s] < \mathit{opt\text{-}val}(s)\} .$$

We select one state $s \in \mathit{sub}(V)$ with the highest optimal value, i.e., s satisfies

$$\mathit{opt\text{-}val}(s) = \max \{\mathit{opt\text{-}val}(s') \mid s' \in \mathit{sub}(V)\} .$$

We show that after (s_k, V_k), which is the first configuration of suffix $\mathcal{X}'$, the value of s strictly increases; this would be a contradiction by choice of $\mathcal{X}'$.

By definition of optimal value, there is an action-path p starting at s, with $\mathit{opt\text{-}val}(s) = \mathit{val}(p)$. Let (s, a) be the first pair of p, and denote $\delta(s, a) = \{s'\}$. By Property 3.18 (below) we execute (s, a) infinitely often in run $\mathcal{X}$.

Property 3.18. In each exploring run, for each $(s, a) \in S \times A$ there are infinitely many transitions in which we execute the pair (s, a).

Proof. Since there are a finite number of configurations (by Lemma 2.7), we can consider a configuration (s_1^*, V_1^*) that occurs infinitely often in the run. By connectedness of the task, there is a path in the state space

$$s_1^* \xrightarrow{a_1^*} s_2^* \dots s_n^* \xrightarrow{a_n^*} s_{n+1}^*,$$

where $n \geq 1$ and $s_{n+1}^* = s$. By the built-in fairness assumption of exploring runs (see Section 2.4.1), we infinitely often follow (a_1^*, s_2^*) from configuration (s_1^*, V_1^*). This results in a configuration (s_2^*, V_2^*) that also occurs infinitely often. The reasoning can be repeated to arrive at a configuration (s_{n+1}^*, V_{n+1}^*) with $s_{n+1}^* = s$ that occurs infinitely often. The reasoning can now be applied one more time. Denoting $\delta(s, a) = \{s'\}$, from configuration (s_{n+1}^*, V_{n+1}^*) we infinitely often follow (a, s'). □

Because the part of $\mathcal{X}$ before $\mathcal{X}'$ is finite, the pair (s, a) is executed infinitely often in $\mathcal{X}'$. So, there are infinitely many transitions in $\mathcal{X}'$ of the following form:

$$(s_h, V_h) \xrightarrow{a_h,\, s_{h+1}} (s_{h+1}, V_{h+1}),$$

where $(s_h, a_h) = (s, a)$. We show below that $V_h(s, a) < V_{h+1}(s, a)$. So, as long as the value of s stays strictly below *opt-val*(s) we can strictly increase the value of (s, a). This always leads to a moment where the value of s in its entirety is strictly increased. This is the sought contradiction.

We are left to show $V_h(s, a) < V_{h+1}(s, a)$. Note that $s_{h+1} = s'$ by determinism. Also, by assumption on the unchanging values in $\mathcal{X}'$, we have (1) $V_h[s] = V[s] <$ *opt-val*(s) and (2) $V_h[s'] = V[s']$. In Algorithm 1, the value change during the above transition is

$$\begin{aligned} d &= \max(V_h[s_{h+1}], R(s, a)) - K - V_h[s] \\ &= \max(V[s'], R(s, a)) - K - V[s]. \end{aligned} \tag{3.2}$$

It suffices to show that $d > 0$. We recall from earlier the action path p starting at s, with *opt-val*$(s) =$ *val*(p). By Property 3.19 (below) we have two cases: either *opt-val*$(s) = [\![R(s, a) - K]\!]$ or *opt-val*$(s) =$ *opt-val*$(s') - K$.

Property 3.19. Let $s \in S$ and let p be an action-path for s with *val*$(p) =$ *opt-val*(s). Let (s, a) be the first pair of p, and denote $\delta(s, a) = \{s'\}$. We have either

- *opt-val*$(s) = [\![R(s, a) - K]\!]$; or,
- *opt-val*$(s) =$ *opt-val*$(s') - K$.

(Proof in Appendix B.3.) □

We consider each case in turn.

First case Suppose *opt-val*$(s) = [\![R(s, a) - K]\!]$. Since $0 \leq V[s] <$ *opt-val*(s), we have *opt-val*$(s) > 0$, so we write more simply *opt-val*$(s) = R(s, a) - K$. Using Equation (3.2),

$$\begin{aligned} d &= \max(V[s'], R(s, a)) - K - V[s] \\ &\geq R(s, a) - K - V[s] \\ &= \textit{opt-val}(s) - V[s]. \end{aligned}$$

Since $V[s] <$ *opt-val*(s) by assumption, we obtain $d > 0$.

Second case Suppose $\textit{opt-val}(s) = \textit{opt-val}(s') - K$. Since $K \geq 1$, we have $\textit{opt-val}(s) < \textit{opt-val}(s')$. Necessarily $s' \notin \textit{sub}(V)$, because otherwise $\textit{opt-val}(s) < \max\{\textit{opt-val}(s'') \mid s'' \in \textit{sub}(V)\}$, which is false by choice of s. Therefore $V[s'] = \textit{opt-val}(s')$. We now complete the reasoning, continuing from Equation (3.2):

$$\begin{aligned} d &= \max(V[s'], R(s,a)) - K - V[s] \\ &\geq V[s'] - K - V[s] \\ &= \textit{opt-val}(s') - K - V[s] \\ &= \textit{opt-val}(s) - V[s]. \end{aligned}$$

Like in the previous case, since $V[s] < \textit{opt-val}(s)$ by assumption, we obtain $d > 0$.

3.4 Proof of Corollary 3.2

Let $T = (S, S^{\text{start}}, A, \delta, R)$ be a DC task. Let $\mathcal{X}$ be an exploring run, and let $\mathcal{X}'$ be the infinite suffix where all value functions are both optimal and consistent, as given by Theorem 3.1. We show that in $\mathcal{X}'$ the value function eventually becomes fixed.

By Property 3.20 (below), we know that for each state the set of preferred actions is fixed throughout $\mathcal{X}'$. For each $s \in S$, let $\textit{pref}(s) \subseteq A$ denote the final set of actions preferred by s in $\mathcal{X}'$. Also by Property 3.20, for each $(s,a) \in S \times A$, if $a \notin \textit{pref}(s)$, we know that the value of the non-preferred pair (s,a) can never be increased in $\mathcal{X}'$. Therefore, the value of non-preferred state-action pairs becomes constant.

Hence, there is a suffix $\mathcal{X}''$ of $\mathcal{X}'$ in which states always prefer the same actions and in which the value of non-preferred state-action pairs is constant. We show that all value functions in $\mathcal{X}''$ are the same. Thereto, let us consider two configurations (s_i, V_i) and (s_j, V_j) in $\mathcal{X}''$. We show for each $(s,a) \in S \times A$ that $V_i(s,a) = V_j(s,a)$. We distinguish between the following cases:

- Suppose $a \notin \textit{pref}(s)$. Then $V_i(s,a) = V_j(s,a)$ by choice of $\mathcal{X}''$.
- Suppose $a \in \textit{pref}(s)$. Then $a \in \textit{pref}(s, V_i)$ and $a \in \textit{pref}(s, V_j)$. Subsequently, $V_i(s,a) = V_i[s]$ and $V_j(s,a) = V_j[s]$. Moreover, by optimality of V_i and V_j,
$$V_i[s] = \textit{opt-val}(s) = V_j[s].$$
Overall, $V_i(s,a) = V_j(s,a)$.

Property 3.20. Consider a transition in $\mathcal{X}'$,

$$(s_i, V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1}, V_{i+1}).$$

We have

1. $\textit{pref}(s_i, V_{i+1}) = \textit{pref}(s_i, V_i)$;
2. If $a_i \notin \textit{pref}(s_i, V_i)$ then $V_{i+1}(s_i, a_i) \leq V_i(s_i, a_i)$.

Note: only the preferred actions of s_i could change; hence, for each $s \neq s_i$ we have $\textit{pref}(s, V_{i+1}) = \textit{pref}(s, V_i)$.

Proof. We distinguish between two cases, depending on whether a_i is preferred or not.

Preferred action If $a_i \in \mathit{pref}(s_i, V_i)$ then consistency of V_i allows us to apply Lemma 3.21 (below) to know $V_{i+1} = V_i$, i.e., executing preferred actions does not modify the value function. Hence, $\mathit{pref}(s_i, V_{i+1}) = \mathit{pref}(s_i, V_i)$.

Non-preferred action Suppose $a_i \notin \mathit{pref}(s_i, V_i)$. By Algorithm 1, we have $V_{i+1}(s_i, a_i) = [\![V_i(s_i, a_i) + d]\!]$ where

$$d = \max(V_i[s_{i+1}], R(s_i, a_i)) - K - V_i[s_i].$$

We show below that $d \leq 0$, which implies $V_{i+1}(s_i, a_i) \leq V_i(s_i, a_i)$. Moreover, $\mathit{pref}(s_i, V_{i+1}) = \mathit{pref}(s_i, V_i)$: since $a_i \notin \mathit{pref}(s_i, V_i)$, the only way to change the set of preferred actions would be to make a_i into a preferred action by a strict value increase (which does not happen).

We are left to show $d \leq 0$. To start, we use that $x \leq [\![x]\!]$ for each integer x; hence,

$$d \leq [\![\max(V_i[s_{i+1}], R(s_i, a_i)) - K]\!] - V_i[s_i].$$

Next, by optimality of V_i,

$$d \leq [\![\max(\mathit{opt\text{-}val}(s_{i+1}), R(s_i, a_i)) - K]\!] - \mathit{opt\text{-}val}(s_i).$$

Subsequently, the clamped part on the right-hand side can be simplified with Lemma 3.14, to obtain

$$d \leq \mathit{opt\text{-}val}(s_i) - \mathit{opt\text{-}val}(s_i) = 0.$$

□

Lemma 3.21. For each deterministic task, for each transition

$$(s, V) \xrightarrow{a, s'} (s', V'),$$

if V is consistent and $a \in \mathit{pref}(s, V)$ then $V' = V$.

Proof. By consistency of V, we have

$$V[s] = [\![\max(V[s'], R(s, a)) - K]\!].$$

Next, we look at Algorithm 1. There are two cases:

- Suppose $\max(V[s'], R(s, a)) - K \geq 0$. Then $V[s] = \max(V[s'], R(s, a)) - K$. Hence, $d = 0$ and surely $V' = V$.
- Suppose $\max(V[s'], R(s, a)) - K < 0$. Then $V[s] = 0$, making $V(s, a) = 0$. Also, $d < 0$, causing

  $$V'(s, a) = [\![V(s, a) + d]\!] = [\![d]\!] = 0 = V(s, a).$$

  Overall, $V' = V$.

□

3.5 Proof of Theorem 3.10

Let $T = (S, S^{\text{start}}, A, \delta, R)$ be the DC task. Consider a greedy run that starts with a value function V that is both optimal and consistent. We recall by Lemma 3.9 that the run is an infinite sequence of value-sprints. Moreover, by Lemma 3.21, since the run is greedy, every configuration in the run uses the value function V.

Consider an arbitrary value-sprint in the run:

$$(s_1, V) \xrightarrow{a_1,\, s_2} \dots \xrightarrow{a_{n-1},\, s_n} (s_n, V) \xrightarrow{a_n,\, s_{n+1}} (s_{n+1}, V),$$

where $n \geq 1$. Here, (s_1, V) is not necessarily the first configuration of the run; it could be anywhere in the run. The corresponding action-path is

$$p = (s_1, a_1), \dots, (s_n, a_n).$$

We show $val(p) = opt\text{-}val(s_1)$. For each $i \in \{1, \dots, n\}$, we define the suffix

$$p_i = (s_i, a_i), \dots, (s_n, a_n).$$

Note that $p_1 = p$. Below, we show by induction on $i = n, \dots, 1$ that $V[s_i] = val(p_i)$. This eventually gives $V[s_1] = val(p_1) = val(p)$. Subsequently, since $V[s_1] = opt\text{-}val(s_1)$ by optimality of V, we obtain $val(p) = opt\text{-}val(s_1)$, as desired.

Base case By definition, $val(p_n) = [\![R(s_n, a_n) - K]\!]$. If $V[s_n] = 0$ then $[\![\max(V[s_{n+1}], R(s_n, a_n)) - K]\!] = 0$ by consistency of V, enforcing $R(s_n, a_n) \leq K$. In that case, $val(p_n) = 0 = V[s_n]$.

Henceforth, we assume $V[s_n] > 0$. The consistency property for s_n may now be written as

$$V[s_n] = \max(V[s_{n+1}], R(s_n, a_n)) - K. \tag{3.3}$$

Necessarily $V[s_{n+1}] < R(s_n, a_n)$; otherwise, Equation (3.3) becomes $V[s_n] = V[s_{n+1}] - K$, which violates condition 2 in the definition of value-sprint (Definition 3.1), namely, $V[s_n] > V[s_{n+1}] - K$. Hence, Equation (3.3) becomes

$$V[s_n] = R(s_n, a_n) - K.$$

Since $V[s_n] > 0$, we may write $V[s_n] = [\![R(s_n, a_n) - K]\!]$. Thus $V[s_n] = val(p_n)$.

Inductive step If $n = 1$ then no inductive step is needed. Henceforth, assume $n \geq 2$. Let $i \in \{2, \dots, n\}$. We assume as induction hypothesis that $V[s_i] = val(p_i)$. By Lemma 3.22 (below), we have

$$val(p_{i-1}) = \max([\![R(s_{i-1}, a_{i-1}) - K]\!], [\![val(p_i) - K]\!]).$$

By subsequently applying the induction hypothesis $V[s_i] = val(p_i)$, we get

$$\begin{aligned} val(p_{i-1}) &= \max([\![R(s_{i-1}, a_{i-1}) - K]\!], [\![V[s_i] - K]\!]) \\ &= [\![\max(V[s_i], R(s_{i-1}, a_{i-1})) - K]\!]. \end{aligned}$$

Since $a_{i-1} \in pref(s_{i-1}, V)$ by greediness of the value-sprint, the last line equals $V[s_{i-1}]$ by consistency. Hence, $V[s_{i-1}] = val(p_{i-1})$.

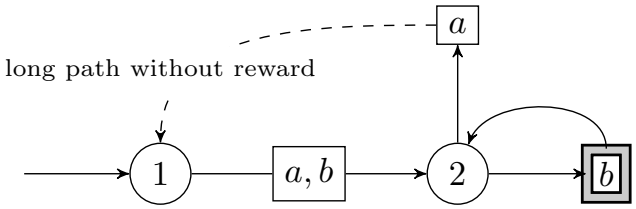

Figure 4.1: A task with start state 1, and at least one other state 2. The available actions are a and b. The only nonzero reward is assigned to $(2, b)$, indicated by a shaded box.

Lemma 3.22. Consider a deterministic task with reward function R. Let $p = (s_1, a_1), (s_2, a_2), \ldots, (s_n, a_n)$ be an action-path with $n \geq 2$. Let $p' = (s_2, a_2), \ldots, (s_n, a_n)$ be the suffix of p after removing the first pair (s_1, a_1). We have

$$val(p) = \max(\llbracket R(s_1, a_1) - K \rrbracket, \llbracket val(p') - K \rrbracket).$$

(Proof in Appendix C.) □

4 Greedy navigation under nondeterminism

As suggested by Example 3.8, optimality is not well-defined for nondeterministic tasks, due to persistent value fluctuations. Therefore, as a measure of agent quality in nondeterministic tasks, we propose to avoid rewardless cycles in the state space. Avoiding cycles is a constraint on the time budget to reach reward. This could be useful, for example in animals, when reward is associated with survival. Exploring runs, however, might repeatedly lead the agent into rewardless cycles. In this section, we show the usefulness of a purely greedy approach to avoid rewardless cycles in nondeterministic navigation problems.

We recall the definition of navigation problems from Definition 3.2 in Section 3.2.1.

4.1 Greedy navigation

We recall from Section 2.4.1 that in a greedy run the agent is constantly following preferred actions, without exploring other possibilities. This corresponds to setting $\epsilon = 0$ in Algorithm 2. We emphasize that a random action is chosen from the preferred actions. This reflects that the agent deems all preferred actions as equally desirable. The agent can only behave purely randomly on a state when it prefers no actions on that state.

The following Example 4.1 motivates the use of greedy runs.

Example 4.1. Consider the task in Figure 4.1. The reward function assigns a nonzero reward only to the pair $(2, b)$. Exploring runs do not try to avoid cycles, and therefore the agent could witness very long cycles without reward if at state 2 the action a is selected several times in succession. This suggests to use greedy runs as a possible way to eventually avoid cycles. □

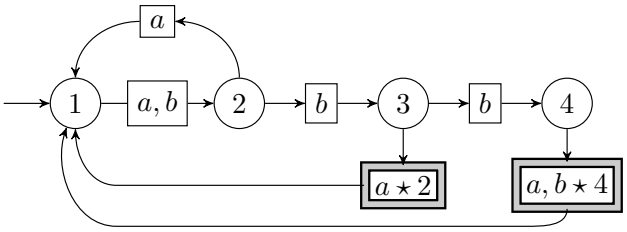

Figure 4.2: Using the same graphical notation as in Figure 4.1, this task has the following nonzero rewards: $(3, a) \mapsto 2$, $(4, a) \mapsto 4$, and $(4, b) \mapsto 4$. The rewarding state-action pairs are indicated by a shaded box, and the reward quantity is indicated by a star.

An important assumption in navigation problems, as defined in Section 3.2.1, is that the reward is large enough to bridge large distances in the state space. The following Example 4.2 illustrates why the greedy approach sometimes fails to avoid cycles when reward is too small. We will therefore restrict attention to navigation problems, where the issue of small reward does not occur.

Example 4.2. We consider the task in Figure 4.2. Suppose for simplicity that $K = 1$. If a greedy run starts with a zero value function, and we would perform $(3, a)$ before $(3, b)$ then the value of $(3, a)$ will become 1. Subsequently, by greediness, action a will be executed whenever the agent visits state 3. Value 1 is however too small to be propagated towards $(2, b)$, and the agent will remain stuck with a value of zero for both $(2, a)$ and $(2, b)$. This way, actions a and b are both preferred in state 2, possibly causing the greedy run to witness long cycles without reward if action a would be chosen at state 2 several times in succession.

Some greedy runs, however, will perform at least twice $(3, b)$ before $(3, a)$. This causes $(3, b)$ to be assigned a value of 2, which in turn causes $(2, b)$ to be assigned value 1. In that scenario, the agent will henceforth never witness cycles without reward. □

4.2 Reducibility

Let $T = (S, S^{\text{start}}, A, \delta, R)$ be a navigation problem, with nonzero reward M. To make assumptions about nondeterminism, we formalize a notion called reducibility, generalizing solvability mentioned for DC navigation problems in Section 3.2.1. We first define

$$\mathit{rewards}(T) = \{(s, a) \in S \times A \mid R(s, a) = M\},$$

and,

$$\mathit{goals}(T) = \{s \in S \mid \exists a \in A \text{ with } (s, a) \in \mathit{rewards}(T)\}.$$

Intuitively, $\mathit{goals}(T)$ contains the states where immediate reward can be obtained. Now, we define the set of reducible states of T as follows,

$$\mathit{reduce}(T) = \bigcup_{i=1}^{\infty} L_i(T),$$

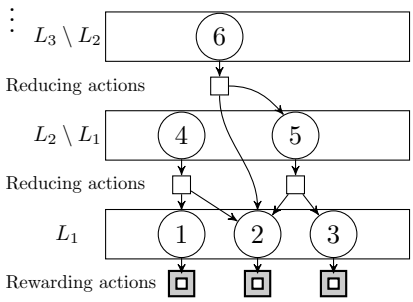

Figure 4.3: Illustration of the layer structure in reducibility.

where

- $L_1(T) = \mathit{goals}(T)$; and,
- for each $i \geq 2$,

$$L_i(T) = L_{i-1}(T) \cup \{s \in S \mid \exists a \in A \text{ with } \delta(s,a) \subseteq L_{i-1}(T)\}.$$

Intuitively, $\mathit{reduce}(T)$ represents a stack of layers, as illustrated in Figure 4.3. Set $L_1(T)$ is the base layer, containing the goal states. Set $L_2(T)$ adds those states that have an action leading into $L_1(T)$, closer to reward. We keep stacking layers until we can add no more states. Despite the nondeterminism in each state-action application, each state in $\mathit{reduce}(T)$ can still approach reward. Since S is finite, there is always a smallest index $n \geq 1$ such that $L_n(T) = L_{n+1}(T)$, i.e., $L_n(T)$ is the fixpoint of the sequence.

We abbreviate $\mathit{non\text{-}reduce}(T) = S \setminus \mathit{reduce}(T)$. Note that for each $s \in \mathit{non\text{-}reduce}(T)$, for each $a \in A$, we must have $\delta(s,a) \cap \mathit{non\text{-}reduce}(T) \neq \emptyset$; otherwise $s \in \mathit{reduce}(T)$. This means that once the agent enters a non-reducible state, the nondeterminism can keep the agent inside the non-reducible states for arbitrary amounts of time. There is no reward in non-reducible states: in a state s, if there would be a rewarding action then $s \in \mathit{goals}(T) \subseteq \mathit{reduce}(T)$.

Definition 4.1. We say that a task $T = (S, S^{\text{start}}, A, \delta, R)$ is reducible if the following conditions are satisfied:

1. $S^{\text{start}} \subseteq \mathit{reduce}(T)$; and,
2. for each $s \in \mathit{non\text{-}reduce}(T)$ and each $s' \in S^{\text{start}}$ there is a path

$$s_1 \xrightarrow{a_1} \dots \xrightarrow{a_{n-1}} s_n,$$

where $(s_1, s_n) = (s, s')$ and $\{s_2, \dots, s_{n-1}\} \subseteq \mathit{non\text{-}reduce}(T)$.[5]

The first condition says that all start states should have a strategy to reward. Whenever the agent would stumble onto a non-reducible state, the second condition provides an escape route back to any start state, entirely tunneled through non-reducible states.

[5]We allow $n = 2$, in which case the path consists of a single jump.

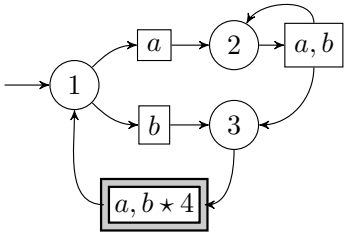

Figure 4.4: Using the same graphical notation as in Figure 4.1, this task has the following nonzero rewards: $(3, a) \mapsto 4$, $(3, b) \mapsto 4$.

To motivate the assumption of reducibility, the following example gives a non-reducible navigation problem that could forever cause cycles without reward, even with greedy runs.

Example 4.3. We consider the task shown in Figure 4.4. For simplicity, we assume $K = 1$ and that initial values are zero. State 2 is the only non-reducible state. When inside state 2, nondeterminism can keep the agent inside state 2 for arbitrary amounts of time, leading to cycles without reward. Value-Ramp can however not always learn to avoid entering state 2. In a greedy run, if the agent would perform $(1, a)$ three times before $(1, b)$ then the following happens: $(3, a)$ or $(3, b)$ gets value 3; next, $(2, a)$ or $(2, b)$ gets value 2; and, $(1, a)$ gets value 1.[6] But then the agent would keep running into state 2; in that case, the greedy run could witness many cycles without reward.

Note that if there would have been an additional escape option from state 2 back to state 1, say $1 \in \delta(2, a)$, then condition 2 of reducibility would be satisfied. In that case, if we are stuck in state 2 long enough, both actions a and b are repeatedly tried, whose value is diminished to zero; that is the right moment to jump back to state 1. The subsequent visit from 1 to 2, through action a, will make the value of $(1, a)$ also zero. To try action b at 1, we should however return from 2 to 1 before witnessing the rewarding exit to state 3. In a greedy run, the built-in fairness condition ensures that this right sequence of events can not be postponed indefinitely. □

Remark 4.4 (Relationship to relocations)**.** For a reducible state, the agent can trust that certain actions will bring the agent (gradually) closer to reward. On reducible states, we rule out that an external observer could intervene at arbitrary moments to send the agent to specific places in the state space. This ensures that the agent can in principle follow nice ramp shapes that peak at reward. By contrast, if we would suddenly relocate the agent without reward to a low-value state, or without reward to a high-value state, then the ramp-shape of the values might be locally damaged. This is suggested in Figure 4.5. □

4.3 Restartability

We also consider an additional technical assumption on tasks:

[6] Recall that greedy runs, as defined in Section 2.4.1, have a built-in fairness condition that would prevent the agent from being stuck inside state 2 forever.

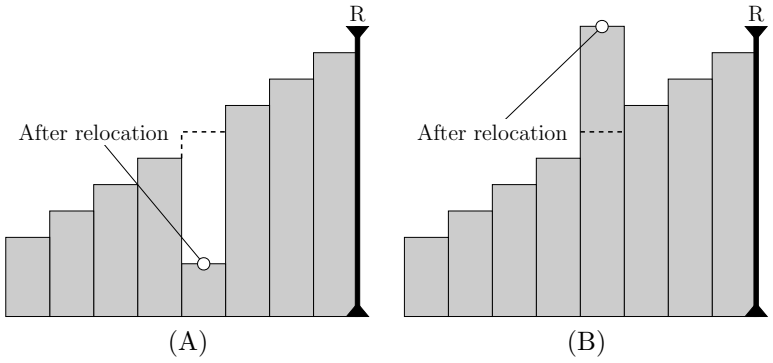

Figure 4.5: When the agent is relocated at wrong moments, an otherwise good ramp shape might be disrupted. The value could be artificially decreased (A) or increased (B).

Definition 4.2. We say that a navigation problem $T = (S, S^{\text{start}}, A, \delta, R)$ is restartable if for each $(s, a) \in \mathit{rewards}(T)$ we have $\delta(s, a) = S^{\text{start}}$.

This assumption allows the environment to put the agent at another start state after reaching a goal. Possibly such start states are very near to the recently obtained reward, making the movement of the agent sometimes appear seamless in the state space. This observation indicates that restartable navigation problems encompass some practical navigation cases on a map, where sometimes we want to simulate that the agent is simply staying at a certain location after obtaining a reward.

In combination with reducibility, the assumption of restartability ensures that we remain inside reducible states once we obtain reward. This way, the agent can in principle continually navigate towards reward without getting trapped in non-reducible states. The following example shows that reducibility by itself does not ensure that the agent can learn to avoid rewardless cycles, but that the combination with restartability is useful.

Example 4.5. We consider the navigation problem shown in Figure 4.6(A). There are two states 1 and 2, and one possible action a. The task is not restartable. After obtaining reward through the pair $(1, a)$, the agent could be trapped inside the non-reducible state 2 for arbitrary amounts of time, leading to rewardless cycles.

In Figure 4.6(B), we consider a modification of subfigure (A) to a reducible and restartable task. Note that state 2 has now become a start state. We have also removed state 2 as a successor state of the pair $(2, a)$, ensuring that 2 is reducible (which is a property demanded for start states by reducibility). In this simple example it is immediately clear that every cycle contains reward. More generally, in Theorem 4.6 (below), we will see the useful effect of combining reducibility and restartability on learning in greedy runs. □

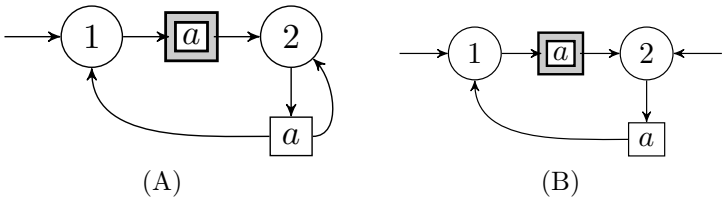

Figure 4.6: (A) reducible but non-restartable navigation problem. (B) reducible and restartable navigation problem. Rewarding actions are indicated by a shaded box.

4.4 Navigation result

As an abbreviation, we say that a navigation problem is RR if the problem is both reducible and restartable. On RR navigation problems, Value-Ramp successfully learns to avoid rewardless cycles in every greedy run:

Theorem 4.6. On each RR navigation problem with reward quantity M, when initial values are below M, in each greedy run, eventually all state cycles contain reward.

The proof is given in Section 4.5.

Remark 4.7 (Assumptions)**.** Removing the assumption on initial values in Theorem 4.6 could be an item for future work. Without the assumption, the agent requires additional time to unlearn high violating values (see Figure 2.3), before it could learn paths towards the true reward. The assumption fortunately does not seem too severe, because a practical simulation might initialize the value function to satisfy the assumption.

Moreover, the notions of reducibility and restartability might perhaps be combined into a more tight concept, where we assume that after obtaining reward we do not necessarily end up at a start state but just at a reducible state. Formally, letting $T = (S, S^{\text{start}}, A, \delta, R)$ be a navigation problem, for each $(s, a) \in \mathit{rewards}(T)$, we could assume $\delta(s, a) \subseteq \mathit{reduce}(T)$. However, not giving a special role to start states might make it more difficult to assume a structure on the non-reducible states; the current assumption in Definition 4.1 is anchored on start states. □

Remark 4.8 (Applicability)**.** Theorem 4.6 works in particular for deterministic RR navigation problems, where, necessarily, there could be only one start state. Theorem 4.6 also applies to RR navigation problems that are completely deterministic on non-rewarding state-action pairs, but where the rewarding state-action pairs are non-deterministic (in selecting the next start state). □

Example 4.9 (RR grid navigation)**.** We extend the simulation of Example 3.3 with nondeterministic effects. We again consider a navigable 2D grid, with actions left, right, up, down, and finish. This time we allow multiple start cells. We keep using the goal cells from earlier: when performing the finish action at goal cells, reward is obtained and the agent is transported back to a randomly selected start cell; this corresponds to the restartability assumption. We add two additional types of cell: swamp cells and jump cells. In a swamp cell, for

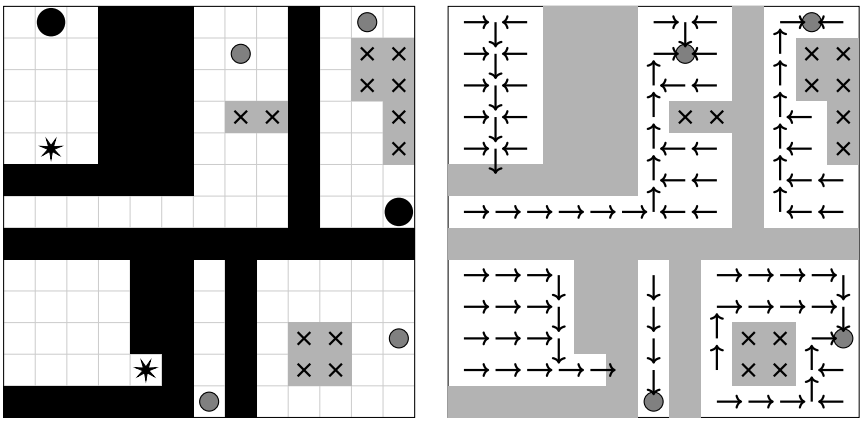

(A) (B)

Figure 4.7: In the context of Example 4.9, we give an RR navigation problem in 2D. (A) We use the following notational convention: start cells → black circle; goal cells → gray circle; swamp cells → "X"; jump cells → star. Reducibility is shown in subfigure (B). Note that the jump cells prefer an action that seemingly jumps through a wall; that is possible because the direction is multiplied by either 2 or 4.

each action, we nondeterministically (1) send the agent back to a (random) start state, or (2) we apply a random offset from the set

$$\{(0,0),(-1,0),(1,0),(0,1),(0,-1)\}.$$

The ability to restart the task from the swap cells is needed for reducibility (condition 2 in Definition 4.1).[7] Due to the offset $(0,0)$, the agent could become stuck for arbitrary amounts of time when it enters a swamp cell. The other offsets let swamp cells unpredictably move the agent; but that is not crucial for this example. Clearly, swamp cells have no action that is guaranteed to reach a goal cell, even if some start cells are goal cells. Hence, all swamp cells are non-reducible.

Second, we have special jump cells: for any movement action, the jump cell takes the direction of the action and nondeterministically multiplies it by either 2 or by 4. For example, on a jump cell, if the agent performs the action 'right', with direction $(1,0)$, then the effective offset could be $(2,0)$ or $(4,0)$. No movement is performed if the resulting offset would end in an obstacle.

We show a concrete example situation in Figure 4.7. Figure 4.8 shows results from one simulation, where the agent eventually avoids rewardless cycles, but where the agent forever fluctuates between different effective paths due to the jump cells. □

4.5 Proof of Theorem 4.6

Let $T = (S, S^{\text{start}}, A, \delta, R)$ be an RR navigation problem, and let M denote the nonzero reward quantity. We fix T throughout this subsection.

[7] Satisfaction of condition 1 in Definition 4.1 depends on the specific 2D map.

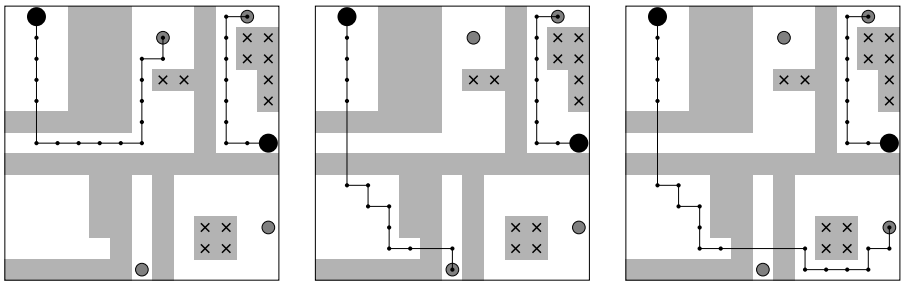

(a) Path 1 (b) Path 2 (c) Path 3

Figure 4.8: Consider the task in Figure 4.7 (A). For the top-left start cell, inside one simulation of a greedy run, the path of the agent to reward strongly depends on the nondeterministic outcome of the jump cells. For the rightmost start cell there is a deterministic path to its local reward.

4.5.1 Proof intuition

The agent repeatedly begins in the start states and finds for each start state some path to reward. The formed paths are not necessarily the shortest. The agent first remembers some high-value actions nearest to reward; such state-action associations form an initial reward strategy, very localized near the reward. Any other states can be gradually added to that initial reward strategy, as suggested in Figure 4.9. There is a growth process of the reward strategy that propagates back to the start states, until the agent has a reward strategy from every start state. At that point, the action preferences strongly restrict what part of the state space is visited by the agent.

We recall that the considered navigation problem is restartable, implying that after obtaining reward we go back to a start state.

We also recall that the navigation problem may be nondeterministic, meaning that the successor state can vary between different applications of the same state-action pair. In general, the agent can not avoid nondeterminism, so the agent should find a reward strategy that steadily moves towards reward despite the nondeterminism; this is illustrated in Figure 4.10. Reducibility (Definition 4.1) ensures the existence of such a reward strategy.

Also, the agent can be distracted by the initial value function, or by early estimated values that are misleading due to nondeterminism. But, eventually, the agent will learn where the reward is. The agent can be thought of as digging through layers of violating values to reach the reward. This requires that wrong values must be modified, either increased or decreased, to match the values emanating from the reward. In the proof, we use the built-in fairness assumption of exploring runs to confront the agent with wrong values.

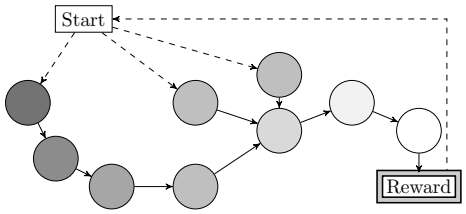

Figure 4.9: The agent gradually builds longer paths to reward by adding new state-action pairs in front of existing paths. In a greedy run, we do not necessarily find the shortest paths to reward. In this figure, states with higher value have a brighter shade.

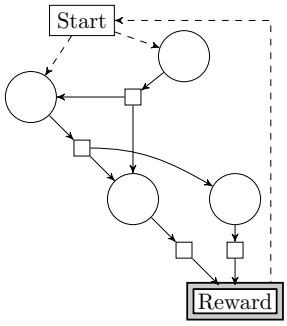

Figure 4.10: Illustration of nondeterminism in the final reward strategy. States are represented by circles, and their chosen actions by boxes. Note that in this figure, each action gets strictly closer to reward, despite the nondeterminism. Reducible tasks always have such a reward strategy.

4.5.2 Strategy

Let V be a value function. Similar to reducibility (Definition 4.1), we collect states that have a reward strategy under V. Formally,

$$\mathit{strategy}(V) = \bigcup_{i=1}^{\infty} z_i(V),$$

where

- the set $z_1(V)$ contains all states s satisfying $V[s] = M - K$, and, $\forall a \in \mathit{pref}(s, V)$ we have
$$(s, a) \in \mathit{rewards}(T).$$
- for each $i \geq 2$, the set $z_i(V)$ extends $z_{i-1}(V)$ with all states s satisfying $V[s] = M - iK$, and, $\forall a \in \mathit{pref}(s, V)$:
  1. $(s, a) \notin \mathit{rewards}(T)$,
  2. $\delta(s, a) \subseteq z_{i-1}(V)$,
  3. $\exists s' \in \delta(s, a)$ with $V[s'] = M - (i - 1)K$.

We call the sets $z_i(V)$ strategy layers. Similar to reducibility, the strategy starts with $z_1(V)$, containing those states that prefer only rewarding actions and have correct value estimation. Then we add layers of states that prefer non-rewarding actions but whose successor states end up closer to $z_1(V)$. Increasing index i corresponds to following value ramps downhill.

We emphasize the following conservative value estimation: for $i \geq 2$, letting $s \in z_i(V)$ and $a \in \mathit{pref}(s, V)$, some successors in $\delta(s, a)$ could have value strictly larger than $M-(i-1)K$, yet we demand $V[s] = M-iK$. This will be important in Section 4.5.4, where we would like strategies to be preserved under some transformation function.

4.5.3 Good configurations

We say that a configuration $c = (s, V)$ is good if the following conditions are satisfied:

1. $S^{\text{start}} \subseteq \mathit{strategy}(V)$; and,
2. $s \in \mathit{strategy}(V)$.

By Property 4.10 (below), we know that once a greedy run encounters a good configuration, all cycles contain reward. The intuition, is that the strategy brings the agent to reward in an acyclic manner. Once all start states and the current state belong to the strategy, the agent is bound inside the strategy forever. In particular, during a reward transition, the agent is sent to another start state, which is inside the strategy.

Our goal is to show that each greedy run eventually encounters a good configuration.

Property 4.10. In every greedy run, after reaching a good configuration, all state cycles contain reward. *(Proof in Appendix D.3.)* □

4.5.4 Transforming configurations

We define a deterministic function β to transform any configuration into a good configuration. More specifically, when given a configuration, function β tells us (1) which action should be taken, and (2) which successor state should be visited. As we will see in Section 4.5.5, the fairness assumption of greedy runs allows β to be called sufficiently often. To specify β as fully deterministic, we assume a total order on the finite sets S and A. Usage of the order is indicated by the "min"-operator on some sets of states and actions.

Restarts from non-reducible states Before defining β, we define an acyclic movement strategy from non-reducible states to start states. Fixing a start state s, we specify a function $go^{(s)} : \mathit{non\text{-}reduce}(T) \to A \times S$ below.

We introduce some convenience notation. Let $X \subseteq S \times A \times S$ be a set. We define $\mathit{dom}\,(X) = \{s \in S \mid \exists a, s' : (s, a, s') \in X\}$. The set X can be viewed as a nondeterministic function $f : \mathit{dom}\,(X) \to \mathcal{P}(A \times S)$, defined for each $s \in \mathit{dom}\,(X)$ as $f(s) = \{(a, s') \mid (s, a, s') \in X\}$.

Now, we define a set $g^{(s)} \subseteq \mathit{non\text{-}reduce}(T) \times A \times S$,

$$g^{(s)} = \bigcup_{i=1}^{\infty} g_i^{(s)},$$

where

- $g_1^{(s)} = \{(s', a, s) \mid s' \in \mathit{non\text{-}reduce}(T), a \in A, s \in \delta(s', a)\}$,
- for each $i \geq 2$,

$$g_i^{(s)} = g_{i-1}^{(s)} \cup \Big\{(s', a, s'') \mid s' \in \mathit{non\text{-}reduce}(T) \setminus \mathit{dom}\left(g_{i-1}^{(s)}\right), \\ a \in A, s'' \in \delta(s', a) \cap \mathit{dom}\left(g_{i-1}^{(s)}\right)\Big\}$$

Intuitively, $g^{(s)}$ specifies how non-reducible states could choose action-successor pairs to move closer to the fixed start state s. Note that this movement solution is acyclic by definition of $g_i^{(s)}$ for each $i \geq 2$. By Property 4.11 (below) we know $\mathit{non\text{-}reduce}(T) \subseteq \mathit{dom}\left(g^{(s)}\right)$.

We convert $g^{(s)}$ into a deterministic function $go^{(s)} : \mathit{non\text{-}reduce}(T) \to A \times S$. Using the assumed order on S and A, for each $s' \in \mathit{non\text{-}reduce}(T)$, we define $go^{(s)}(s')$ as the lexicographically smallest pair in the set $g^{(s)}(s')$.

Property 4.11. For each start state s, we have $\mathit{non\text{-}reduce}(T) \subseteq \mathit{dom}\left(g^{(s)}\right)$. *(Proof in Appendix D.5.)* □

The function β Let $c = (s, V)$ be an input configuration for β.

To jump outside $\mathit{strategy}(V)$ during reward, we define

$$\mathit{restart}(c) = \begin{cases} \min(S^{\text{start}} \setminus \mathit{strategy}(V)) & \text{if } S^{\text{start}} \not\subseteq \mathit{strategy}(V) \\ \min(S^{\text{start}}) & \text{otherwise.} \end{cases}$$

Function β produces an action-successor pair based on the following nested case analysis. The main idea is to gradually decrease wrong values, to arrive at value zero; at that moment we can pull the agent into any desirable direction, in particular towards reward.

1 Suppose $V[s] = 0$. Then all actions are equally preferable in s. We choose an action and successor to gradually bring us closer to reward. There are three mutually disjoint cases.

1.1 Suppose $s \in \mathit{non\text{-}reduce}(T)$.[8] Let $s_0 = \mathit{restart}(c)$.
$\Rightarrow$ Define $\beta(c) = \mathit{go}^{(s_0)}(s)$.

1.2 Suppose $s \in \mathit{goals}(T)$.[9] Let $a = \min\{a' \in A \mid (s, a') \in \mathit{rewards}(T)\}$ and $s_0 = \mathit{restart}(c)$.[10]
$\Rightarrow$ Define $\beta(c) = (a, s_0)$.

1.3 Suppose $s \in \mathit{reduce}(T) \setminus \mathit{goals}(T)$. We move one layer down into the reducibility structure. Recalling $\mathit{reduce}(T) = \bigcup_{i=1}^{\infty} L_i(T)$ from Section 4.2, we write $\mathit{layer}(s)$ to denote the smallest index i for which $s \in L_i(T)$. Let $a = \min\left\{a' \in A \mid \delta(s, a') \subseteq L_{\mathit{layer}(s)-1}(T)\right\}$, and

$$s' = \begin{cases} \min(\delta(s,a) \setminus \mathit{strategy}(V)) & \text{if } \delta(s,a) \not\subseteq \mathit{strategy}(V) \\ \min(\delta(s,a)) & \text{otherwise.} \end{cases}$$

$\Rightarrow$ Define $\beta(c) = (a, s')$.

2 Otherwise $V[s] > 0$. We confront the agent with wrong values, if any, that lead away from reward. Importantly, because the agent is always greedy in the context of this proof, we may only choose actions preferred by the agent, i.e., actions from $\mathit{pref}(s, V)$.

2.1 Suppose there is an action $a \in \mathit{pref}(s, V)$ for which $(s, a) \notin \mathit{rewards}(T)$ and $\delta(s, a) \not\subseteq \mathit{strategy}(V)$. Let a be the smallest from such actions. Let $s' = \min(\delta(s, a) \setminus \mathit{strategy}(V))$.
$\Rightarrow$ Define $\beta(c) = (a, s')$.

2.2 Otherwise, for all $a \in \mathit{pref}(s, V)$ we have either $(s, a) \in \mathit{rewards}(T)$ or $\delta(s, a) \subseteq \mathit{strategy}(V)$. There could still be errors in the value estimation. Regarding notation, for any $a \in A$ with $\delta(s, a) \subseteq \mathit{strategy}(V)$, we define

$$\mathit{expect}(V, s, a) = \min\{V[s'] \mid s' \in \delta(s, a)\}.$$

Intuitively, $\mathit{expect}(V, s, a)$ is a conservative value expectation.

2.2.1 Suppose there is an action $a \in \mathit{pref}(s, V)$ with $(s, a) \in \mathit{rewards}(T)$ but $V(s, a) \neq M - K$. Let a be the smallest such action, and let $s_0 = \mathit{restart}(c)$.
$\Rightarrow$ Define $\beta(c) = (a, s_0)$.

2.2.2 Suppose there is an action $a \in \mathit{pref}(s, V)$ with $(s, a) \notin \mathit{rewards}(T)$, and therefore $\delta(s, a) \subseteq \mathit{strategy}(V)$, with

$$V(s, a) \neq \mathit{expect}(V, s, a) - K.$$

Let a be the smallest such action, and let

$$s' = \min\{s'' \in \delta(s, a) \mid V[s''] = \mathit{expect}(V, s, a)\}.$$

$\Rightarrow$ Define $\beta(c) = (a, s')$.

[8] We have $s \notin S^{\text{start}}$ because $S^{\text{start}} \subseteq \mathit{reduce}(T)$ by assumption (Definition 4.1).
[9] Recall that $\mathit{goals}(T) \subseteq \mathit{reduce}(T)$ (Section 4.2).
[10] The restartability assumption on the task tells us that $s_0 \in \delta(s, a)$.

2.2.3 Otherwise the value estimation is correct. We choose an action-successor pair to proceed. Let $a = \min(\mathit{pref}(s, V))$.

2.2.3.1 Suppose $(s, a) \in \mathit{rewards}(T)$. Let $s_0 = \mathit{restart}(c)$.
$\Rightarrow$ Define $\beta(c) = (a, s_0)$.

2.2.3.2 Otherwise, $(s, a) \notin \mathit{rewards}(T)$, but we still know $\delta(s, a) \subseteq \mathit{strategy}(V)$. Let

$$s' = \min \{s'' \in \delta(s, a) \mid V[s''] = \mathit{expect}(V, s, a)\}.$$

$\Rightarrow$ Define $\beta(c) = (a, s')$.

4.5.5 Eventually good configurations

We fix a greedy run. By means of function β, we show that the greedy run eventually encounters a good configuration. Intuitively, β represents the useful learning opportunities that are witnessed by the agent.

Bring start states in strategy First, we show the existence of a configuration (s, V) that occurs infinitely often in the run and with $S^{\text{start}} \subseteq \mathit{strategy}(V)$.

Because there are only a finite number of configurations (Lemma 2.7), there is at least one configuration $c = (s, V)$ that occurs infinitely often. Note that function β, by design, proposes action-successor pairs allowed by a greedy transition. Since c occurs infinitely often, the fairness assumption of greedy runs (see Section 2.4.1) tells us that we perform the following transition infinitely often:

$$c \xrightarrow{a,s'} (s', V'),$$

where $(a, s') = \beta(c)$. So, (s', V') too occurs infinitely often, and therefore we can also apply β to (s', V'), and so on. We see that β can be applied arbitrarily many times; this process does not necessarily happen as a contiguous sequence of transitions in the run. If $S^{\text{start}} \not\subseteq \mathit{strategy}(V')$ then Property 4.12 (below) tells us that we eventually discover a configuration (s'', V'') with $\mathit{strategy}(V') \subsetneq \mathit{strategy}(V'')$, i.e., with the strategy strictly extended, that occurs infinitely often. As long as the configurations encountered by β have a start state outside the strategy, we can repeat Property 4.12 to strictly extend the strategy. But the strategy can not keep growing because there are a finite number of states. By repeated application of β, we eventually arrive at a configuration (s'', V'') with $S^{\text{start}} \subseteq \mathit{strategy}(V'')$ that occurs infinitely often in the greedy run.

Property 4.12. Beginning at a configuration (s, V) with $S^{\text{start}} \not\subseteq \mathit{strategy}(V)$, by repeatedly applying β we eventually reach a configuration (s', V') with $\mathit{strategy}(V) \subsetneq \mathit{strategy}(V')$, i.e., we have strictly extended the strategy. *(Proof in Appendix D.6.)* □

Bring current state in strategy At this point, we have shown that there is a configuration $c = (s, V)$ occurring infinitely often in the run and with $S^{\text{start}} \subseteq \mathit{strategy}(V)$. We proceed to showing the existence of a good configuration, where additionally the current state is in the strategy. Again, by the fairness assumption on the run, we can apply β an arbitrary number of times, starting at configuration c. By Property 4.13 (below) we know that the strategy is preserved. Moreover, by Property 4.14 (below) there is at least one occurrence

of reward. Since the task is restartable, during the reward transition we arrive at a start state, which is inside the strategy. At that moment, we have reached a good configuration, as desired.

Property 4.13. Function β always preserves the strategy. More formally, for each transition generated by β,

$$(s_i, V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1}, V_{i+1}),$$

we have *strategy*$(V_i) \subseteq$ *strategy*(V_{i+1}). *(Proof in Appendix D.7.)* □

Property 4.14. Beginning at any configuration, by repeatedly applying β we encounter infinitely many reward transitions. *(Proof in Appendix D.8.)* □

5 Conclusion and further work

By means of formal theorems, we have given concrete insights into the operation of Value-Ramp on well-defined classes of tasks. We now discuss interesting items for further work.

Practical case studies In this paper we have been occupied with the search for general yet nontrivial descriptions of the agent behavior generated by Value-Ramp. A complementary study could focus on testing Value-Ramp on various practical problems, to observe agent behavior on more concrete circumstances, and meanwhile to judge the practical viability of the technique. It appears likely that Value-Ramp can be used for much more problems than the 2D grid examples that we have given. For example, each state could be a sequence of sensory cues, to represent agent conceptualization in a complex environment (Mnih et al., 2015).

In the usage of Value-Ramp, a concrete proposal could be to set the reward quantities rather high and to take $K = 1$, because then the ramps are longer and the agent can subsequently learn long strategies to rewarding events. Moreover, our intuition from the proof of Theorem 4.6 is that nondeterministic tasks could in general be very slow to learn, because the formation of an acyclic rewarding strategy seems to require rare learning opportunities to be (eventually) witnessed. Practical studies might therefore benefit from introducing sufficiently specific concepts inside the agent, so that tasks are rendered approximately deterministic. As suggested by Frémaux et al. (2013), specific concepts might correspond to place cells in the brain, see e.g. (Moser et al., 2008).

Generalized exploration property Perhaps Theorem 3.1 and Theorem 3.10 can be generalized to particular kinds of nondeterministic tasks. Likely, in such a generalization, we should not seek numerical stability of the values, but rather a stability of the knowledge of the highest value paths. This can be likened to Theorem 4.6, of greedy navigation on nondeterministic tasks, where we sought a behavioral stability property instead of a numerical (value) stability property. Of course, it could be that, even on simple tasks, continued exploration leads to continued fluctuations in agent behavior, as suggested by the simple Example 3.8.

More navigation problems Possibly, Value-Ramp can learn to avoid rewardless cycles on more navigation problems than the RR ones of Section 4. More work is needed to understand the form of learnable navigation problems. Some further suggestions on relaxing Theorem 4.6 are mentioned in Remark 4.7.

Negative reward and avoidance Reward in this paper is always a non-negative quantity. Negative quantities could be introduced to study avoidance learning. Or, one could consider a dual value-ramp principle for estimating the aversiveness of state-action pairs. In the aversive value-ramp, the values increase as the agent approaches an aversive stimulus. Whereas greediness in a rewarding value-ramp selects actions to maximize value, greediness in an aversive value-ramp selects actions to minimize value.

Partial observability and features Towards better understanding Value-Ramp on more practical problems, it might be useful to formalize how the task structure is derived from various practical constraints. For example, the agent might have sensors with limited range, leading to perceived states that deviate from the true environment states. This leads to structural assumptions on the transition function δ. It appears interesting to make concrete insights similar to the ones we have presented when more structural assumptions about the tasks are taken into account.

A brain consists of multiple neurons, and each neuron might represent a feature, i.e., a piece of state information. Each encountered task state is projected to a set of features. It appears interesting to extend our framework to learn value for feature-action pairs instead of state-action pairs. In each state, the feature-action pair with the highest value could determine the action for the state.

Relationship with reward discounting Many algorithms in reinforcement learning are based on reward discounting (Sutton and Barto, 1998). An important observation is that reward discounting is based on multiplying values with a rational number γ between zero and one, whereas the value ramp is based on subtracting a strictly positive constant. In further work, it could be interesting to clarify the relationship between reward discounting and the value-ramp principle. The notions could be complementary, but they could be equivalent on certain classes of tasks and reward definitions.

Biological plausibility The Value-Ramp algorithm is inspired by simulations of biologically plausible learning models (Frémaux et al., 2013), that could correspond to observations in biology (van der Meer and Redish, 2011). Possibly, further work could elicit whether suitable variations of Value-Ramp accurately model biological learning. In the current Value-Ramp algorithm, negative updates to value utilize an arbitrary range, i.e., the d-value in Algorithm 1 has no constraints (in particular for the negative range). That might not be biologically realistic: if biological learning is based on dopamine, the negative value updates are likely caused by suppressing dopamine; but the dopamine baseline (in neutral circumstances) is already relative low (Schultz, 2013). One might suspect that multiple iterations of dopamine suppression are needed to unlearn

wrong value expectations. To obtain this effect, we could redefine the learning rule of Equation (2.1) to a rule of the following kind:

$$\Delta_K(v, v', r) = [\![\max(v', r) - v]\!] - K.$$

The effect is that at most K is subtracted when value expectation is not met by successor value or by reward. This could model a limited but noticeable erosion effect on value, in particular on neuronal connections (representing value) during periods of dopamine suppression.

One concrete hypothesis could be that on a ramp-like value experience, the steps of size K represent small dopamine releases that sustain useful concept-action connections in the brain. In absence of such a dopamine release, there could be a net erosion effect on the recently triggered neuronal connections, to unlearn wrong actions.

References

Cormen, T., Leiserson, C., Rivest, R., and Stein, C. (2009). *Introduction to Algorithms, Third Edition*. The MIT Press.

Frémaux, N., Sprekeler, H., and Gerstner, W. (2013). Reinforcement learning using a continuous time actor-critic framework with spiking neurons. *PLoS Computational Biology*, 9(4):e1003024.

Gustafson, J. (2015). *The End of Error: Unum Computing*. Chapman and Hall/CRC.

Hennigan, K., D'Ardenne, K., and McClure, S. (2015). Distinct midbrain and habenula pathways are involved in processing aversive events in humans. *The Journal of Neuroscience*, 35(1):198–208.

Mnih, V., Kavukcuoglu, K., Silver, D., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540):529 – 533.

Moser, E., Kropff, E., and Moser, M. (2008). Place cells, grid cells, and the brain's spatial representation system. *Annual Review of Neuroscience*, 31:69–89.

Potjans, W., Diesmann, M., and Morrison, A. (2011). An imperfect dopaminergic error signal can drive temporal-difference learning. *PLoS Computational Biology*, 7(5):e1001133.

Schultz, W. (2013). Updating dopamine reward signals. *Current Opinion in Neurobiology*, 23(2):229 – 238.

Schultz, W. (2015). Neuronal reward and decision signals: From theories to data. *Physiological Reviews*, 95(3):853–951.

Schwartz, A. (1993). A reinforcement learning method for maximizing undiscounted rewards. In *Machine Learning, Proceedings of the Tenth International Conference*, pages 298–305.

Sutton, R. (1988). Learning to predict by the methods of temporal differences. *Machine Learning*, 3(1):9–44.

Sutton, R. and Barto, A. (1998). *Reinforcement Learning, An Introduction*. The MIT Press.

van der Meer, M. and Redish, A. (2011). Theta phase precession in rat ventral striatum links place and reward information. *The Journal of Neuroscience*, 31(8):2843–2854.

Watkins, C. (1989). Learning from delayed rewards. *PhD thesis, Cambridge University*.

Watkins, C. and Dayan, P. (1992). Q-learning. *Machine Learning*, 8(3–4):279–292.

Appendix

A Proof details of Lemma 2.7

Lemma A.1. For each transition $(s, V) \xrightarrow{a,s'} (s', V')$, we have $\mathit{ceiling}(V) \geq \mathit{ceiling}(V')$.

Proof. The reward quantities never change, and therefore $R(s'', a'') \leq \mathit{ceiling}(V)$ for each $(s'', a'') \in S \times A$. Regarding values, only the value of (s, a) can change during the transition. Therefore $V'(s'', a'') = V(s'', a'') \leq \mathit{ceiling}(V)$ for each $(s'', a'') \in S \times A \setminus \{(s, a)\}$.

There are two cases for (s, a):

- Suppose $V'(s, a) \leq V(s, a)$. Since always $V(s, a) \leq \mathit{ceiling}(V)$, we have $V'(s, a) \leq \mathit{ceiling}(V)$.
- Suppose $V'(s, a) > V(s, a)$. By Algorithm 1, we have

$$\begin{aligned} V'(s,a) &= \llbracket V(s,a) + \big(\max(V[s'], R(s,a)) - K - V[s]\big) \rrbracket \\ &\leq \llbracket V[s] + \big(\max(V[s'], R(s,a)) - K - V[s]\big) \rrbracket \\ &= \llbracket \max(V[s'], R(s,a)) - K \rrbracket \\ &\leq \llbracket \max(\mathit{ceiling}(V), \mathit{ceiling}(V)) - K \rrbracket \\ &\leq \llbracket \mathit{ceiling}(V) - K \rrbracket \\ &\leq \llbracket \mathit{ceiling}(V) \rrbracket \\ &= \mathit{ceiling}(V). \end{aligned}$$

  In the last step we use that always $\mathit{ceiling}(V) \geq 0$. Overall, $V'(s, a) \leq \mathit{ceiling}(V)$.

□

B Proof details of Theorem 3.1

B.1 Proof of Property 3.15

Let V be a value function. Letting $(s, a) \in S \times A$, and denoting $\delta(s, a) = \{s'\}$, we say that (s, a) is a violation in V if

$$V(s, a) > \llbracket \max(V[s'], R(s, a)) - K \rrbracket .$$

We define the highest violation value in V, denoted $\mathit{viol\text{-}max}(V)$, as follows:

$$\mathit{viol\text{-}max}(V) = \begin{cases} \max \{V(s, a) \mid (s, a) \in \mathit{viol}(V)\} & \text{if } \mathit{viol}(V) \neq \emptyset \\ 0 & \text{otherwise.} \end{cases}$$

Always $\mathit{viol\text{-}max}(V) \geq 0$. Also note that $\mathit{viol}(V) = \emptyset \iff \mathit{viol\text{-}max}(V) = 0$:

- If $\mathit{viol}(V) = \emptyset$ then $\mathit{viol\text{-}max}(V) = 0$ by definition.
- Suppose $\mathit{viol}(V) \neq \emptyset$. Each violation (s, a) in V satisfies $V(s, a) > \llbracket \max(V[s'], R(s, a)) - K \rrbracket \geq 0$, implying $\mathit{viol\text{-}max}(V) > 0$.

The following property will be useful:

Property B.1. For each transition $(s_i, V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1}, V_{i+1})$ on T we have

$$\textit{viol-max}(V_i) \geq \textit{viol-max}(V_{i+1}).$$

(Proof in Appendix B.1.1.) □

Recall that the exploring run $\mathcal{X}$ is denoted as

$$(s_1, V_1) \xrightarrow{a_1,\, s_2} (s_2, V_2) \xrightarrow{a_2,\, s_3} \dots$$

We gradually remove all violations. As long as there are violations in $\mathcal{X}$, the highest violation value is strictly positive. So, while there are violations, Property B.2 (below) tells us that the highest violation value can be strictly decreased. There can only be a finite number of such strict decrements because values are at least zero. Hence, eventually the highest violation value becomes zero. Thereafter, all value functions are valid, because for each transition

$$(s_i, V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1}, V_{i+1}),$$

if V_i is valid then $\textit{viol-max}(V_i) = 0$ and therefore $\textit{viol-max}(V_{i+1}) = 0$ by Property B.1, implying $\textit{viol}(V_{i+1}) = \emptyset$.

Property B.2. For each configuration index $i \geq 1$, if $\textit{viol-max}(V_i) > 0$ then there is a configuration index $j > i$ with

$$\textit{viol-max}(V_i) > \textit{viol-max}(V_j),$$

i.e., the highest violation value has been strictly decreased. *(Proof in Appendix B.1.2.)* □

B.1.1 Proof of Property B.1

Consider a transition,

$$(s_i, V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1}, V_{i+1}).$$

We show $\textit{viol-max}(V_i) \geq \textit{viol-max}(V_{i+1})$. To start, by Algorithm 1, we have

$$V_{i+1}(s_i, a_i) = \left[\!\left[V_i(s_i, a_i) + d \right]\!\right],$$

where $d = \max(V_i[s_{i+1}], R(s_i, a_i)) - K - V_i[s_i]$.

Let $(s, a) \in \textit{viol}(V_{i+1})$, and denote $\delta(s, a) = \{s'\}$. We show $V_{i+1}(s, a) \leq \textit{viol-max}(V_i)$; overall, this implies $\textit{viol-max}(V_{i+1}) \leq \textit{viol-max}(V_i)$. We distinguish between the following cases: $d \geq 0$ and $d < 0$.

First case ($d \geq 0$) If $d \geq 0$ then values are not decreased during the transition, implying $V_{i+1}[s'] \geq V_i[s']$. If $(s, a) \neq (s_i, a_i)$ then, using $(s, a) \in \textit{viol}(V_{i+1})$, we have

$$\begin{aligned} V_i(s, a) &= V_{i+1}(s, a) \\ &> \left[\!\left[\max(V_{i+1}[s'], R(s, a)) - K \right]\!\right] \\ &\geq \left[\!\left[\max(V_i[s'], R(s, a)) - K \right]\!\right], \end{aligned}$$

which implies $(s,a) \in \mathit{viol}(V_i)$, and therefore $V_{i+1}(s,a) = V_i(s,a) \leq \mathit{viol\text{-}max}(V_i)$.

We show that the other case, $(s,a) = (s_i,a_i)$, is impossible. Indeed, if $(s,a) = (s_i,a_i)$ then, based on the above equation for $V_{i+1}(s_i,a_i)$, and using $V_i(s_i,a_i) \leq V_i[s_i]$ (which is always true),

$$\begin{aligned} V_{i+1}(s_i,a_i) &= [\![V_i(s_i,a_i) + d]\!] \\ &\leq [\![V_i[s_i] + d]\!] \\ &= [\![\max(V_i[s_{i+1}], R(s_i,a_i)) - K]\!] \\ &\leq [\![\max(V_{i+1}[s_{i+1}], R(s_i,a_i)) - K]\!], \end{aligned}$$

which implies $(s_i,a_i) \notin \mathit{viol}(V_{i+1})$.

Second case ($d < 0$) Note that $d < 0$ implies $V_{i+1}(s,a) \leq V_i(s,a)$.

First, if $(s,a) \in \mathit{viol}(V_i)$ then

$$V_{i+1}(s,a) \leq V_i(s,a) \leq \mathit{viol\text{-}max}(V_i).$$

Henceforth, suppose $(s,a) \notin \mathit{viol}(V_i)$, i.e., $(s,a) \in \mathit{viol}(V_{i+1})$ is a violation newly created during $(s_i,V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1},V_{i+1})$. We observe

$$\begin{aligned} [\![\max(V_{i+1}[s'], R(s,a)) - K]\!] &< V_{i+1}(s,a) \\ &\leq V_i(s,a) \\ &\leq [\![\max(V_i[s'], R(s,a)) - K]\!]. \end{aligned} \tag{B.1}$$

Hence, $V_{i+1}[s'] < V_i[s']$.[11] Therefore $s' = s_i$.

Subsequently, we have $a_i \in \mathit{pref}(s',V_i)$; otherwise there would be some $a' \in \mathit{pref}(s',V_i)$ with $a_i \neq a'$ and $V_{i+1}(s',a') = V_i(s',a') = V_i[s']$, implying $V_i[s'] \leq V_{i+1}[s']$, which is false.

Hence, $V_i(s',a_i) = V_i[s']$. Now, using the definition of d above, the inequality $d < 0$ implies

$$V_i[s'] = V_i[s_i] > \max(V_i[s_{i+1}], R(s_i,a_i)) - K.$$

Combined,

$$V_i(s',a_i) > \max(V_i[s_{i+1}], R(s_i,a_i)) - K,$$

which implies $(s',a_i) \in \mathit{viol}(V_i)$. Thus $V_i(s',a_i) \leq \mathit{viol\text{-}max}(V_i)$.

Lastly, we have $[\![\mathit{viol\text{-}max}(V_i) - K]\!] > [\![R(s,a) - K]\!]$. Otherwise, when considering $[\![\mathit{viol\text{-}max}(V_i) - K]\!] \leq [\![R(s,a) - K]\!]$, Equation (B.1) would imply the following contradiction, using $V_i[s'] = V_i(s',a_i) \leq \mathit{viol\text{-}max}(V_i)$ (from above):

$$\begin{aligned} [\![R(s,a) - K]\!] \leq [\![\max(V_{i+1}[s'], R(s,a)) - K]\!] &< [\![\max(V_i[s'], R(s,a)) - K]\!] \\ &= [\![\max(V_i(s',a_i), R(s,a)) - K]\!] \\ &\leq [\![\max(\mathit{viol\text{-}max}(V_i), R(s,a)) - K]\!] \\ &= \max([\![\mathit{viol\text{-}max}(V_i) - K]\!], [\![R(s,a) - K]\!]) \\ &\leq [\![R(s,a) - K]\!]. \end{aligned}$$

[11] Otherwise, if $V_{i+1}[s'] \geq V_i[s']$ then actually $V_{i+1}[s'] = V_i[s']$ by $d < 0$; subsequently $[\![\max(V_{i+1}[s'], R(s,a)) - K]\!] = [\![\max(V_i[s'], R(s,a)) - K]\!]$, which is false.

So, if $[\![\mathit{viol\text{-}max}(V_i) - K]\!] > [\![R(s,a) - K]\!]$ then, using $(s,a) \notin \mathit{viol}(V_i)$, we observe

$$\begin{aligned} V_{i+1}(s,a) &\leq V_i(s,a) \\ &\leq [\![\max(V_i[s'], R(s,a)) - K]\!] \\ &= [\![\max(V_i(s',a_i), R(s,a)) - K]\!] \\ &\leq [\![\max(\mathit{viol\text{-}max}(V_i), R(s,a)) - K]\!] \\ &\leq [\![\mathit{viol\text{-}max}(V_i) - K]\!]\,. \end{aligned}$$

Moreover, $[\![\mathit{viol\text{-}max}(V_i) - K]\!] \leq \mathit{viol\text{-}max}(V_i)$.[12] Everything combined, we have $V_{i+1}(s,a) \leq \mathit{viol\text{-}max}(V_i)$, as desired.

B.1.2 Proof of Property B.2

Let i be a configuration index, denoting the corresponding configuration as (s_i, V_i), where $\mathit{viol\text{-}max}(V_i) > 0$. By Property B.1, we know for all subsequent configuration indices j with $j \geq i$ that $\mathit{viol\text{-}max}(V_i) \geq \mathit{viol\text{-}max}(V_j)$, i.e., the highest violation value never increases.

Towards a contradiction, suppose that $\mathit{viol\text{-}max}(V_i) = \mathit{viol\text{-}max}(V_j)$ for all $j \geq i$. Because there are only a finite number of configurations by Lemma 2.7, there must be a configuration (s^*, V^*) that occurs infinitely often, and with $\mathit{viol\text{-}max}(V^*) = \mathit{viol\text{-}max}(V_i)$. Since $\mathit{viol\text{-}max}(V_i) > 0$, we can consider a violation $(s,a) \in \mathit{viol}(V^*)$ with $V^*(s,a) = \mathit{viol\text{-}max}(V^*)$. By Property 3.18, there are infinitely many transitions where we execute the pair (s,a).

To continue with the proof, since (s^*, V^*) occurs infinitely often, and (s,a) is infinitely often executed, after configuration (s_i, V_i) we can consider a finite run-fragment F of the following form:

$$(s^*, V^*) \to \ldots \text{ in between execute } (s,a) \text{ at least once} \ldots \to (s^*, V^*).$$

In the fragment F, there must be a *last* transition in which we execute (s,a), denoted as

$$(s_j, V_j) \xrightarrow{a_j,\, s_{j+1}} (s_{j+1}, V_{j+1}),$$

where $(s_j, a_j) = (s,a)$. Since this transition is the last transition of (s,a) in fragment F, we must have $V_{j+1}(s,a) = V^*(s,a) = \mathit{viol\text{-}max}(V^*)$.

There are two cases: either $(s,a) \in \mathit{viol}(V_{j+1})$ or $(s,a) \notin \mathit{viol}(V_{j+1})$. In each case, we derive a contradiction. Denote $\delta(s,a) = \{s'\}$, and abbreviate $W = \mathit{viol\text{-}max}(V^*)$. By assumption at the beginning of this proof, $W > 0$.

First case Suppose $(s,a) \in \mathit{viol}(V_{j+1})$. By Algorithm 1, we have

$$V_{j+1}(s,a) = [\![V_j(s,a) + d]\!]\,,$$

where $d = \max(V_j[s'], R(s,a)) - K - V_j[s]$. It must be $d < 0$. Otherwise, considering $d \geq 0$, since always $V_j(s,a) \leq V_j[s]$, and additionally $V_j[s'] \leq$

[12] If $\mathit{viol\text{-}max}(V_i) - K < 0$ then $[\![\mathit{viol\text{-}max}(V_i) - K]\!] = 0 \leq \mathit{viol\text{-}max}(V_i)$. If $\mathit{viol\text{-}max}(V_i) - K \geq 0$ then $[\![\mathit{viol\text{-}max}(V_i) - K]\!] = \mathit{viol\text{-}max}(V_i) - K < \mathit{viol\text{-}max}(V_i)$.

$V_{j+1}[s']$ when $d \geq 0$, we would have

$$\begin{aligned} V_{j+1}(s,a) &\leq [\![V_j[s] + d]\!] \\ &= [\![\max(V_j[s'], R(s,a)) - K]\!] \\ &\leq [\![\max(V_{j+1}[s'], R(s,a)) - K]\!], \end{aligned}$$

and therefore $(s,a) \notin \mathit{viol}(V_{j+1})$, which is false by assumption.

Note that $V_j(s,a) > V_{j+1}(s,a)$, since $d < 0$ and $V_{j+1}(s,a) = W > 0$. Therefore $V_j(s,a) > W$. We will show below that $(s,a) \in \mathit{viol}(V_j)$, giving

$$\mathit{viol\text{-}max}(V_j) \geq V_j(s,a) > W = \mathit{viol\text{-}max}(V^*) = \mathit{viol\text{-}max}(V_i),$$

in particular, $\mathit{viol\text{-}max}(V_j) > \mathit{viol\text{-}max}(V_i)$, which contradicts Property B.1 (since $i \leq j$).

To show $(s,a) \in \mathit{viol}(V_j)$, we consider the following cases.

- Suppose $a \in \mathit{pref}(s, V_j)$. Hence, $V_j[s] = V_j(s,a)$, which we substitute into the equation of $V_{j+1}(s,a)$ given by Algorithm 1:

  $$\begin{aligned} V_{j+1}(s,a) &= [\![V_j(s,a) + \max(V_j[s'], R(s,a)) - K - V_j(s,a)]\!] \\ &= [\![\max(V_j[s'], R(s,a)) - K]\!]. \end{aligned}$$

  Combined with $V_j(s,a) > V_{j+1}(s,a)$ (see above), we obtain

  $$V_j(s,a) > [\![\max(V_j[s'], R(s,a)) - K]\!],$$

  and therefore $(s,a) \in \mathit{viol}(V_j)$.

- Suppose $a \notin \mathit{pref}(s, V_j)$ and $s' = s$. If $a \notin \mathit{pref}(s, V_j)$ then there is some $a' \in \mathit{pref}(s, V_j)$ with $a \neq a'$. Note that $V_{j+1}(s,a') = V_j(s,a') = V_j[s]$, implying $V_j[s] \leq V_{j+1}[s]$. Moreover, $d < 0$ implies $V_{j+1}[s] \leq V_j[s]$. Overall, $V_{j+1}[s] = V_j[s]$.

  Next, since $(s,a) \in \mathit{viol}(V_{j+1})$, we have, substituting $V_{j+1}[s'] = V_{j+1}[s] = V_j[s]$,

  $$\begin{aligned} V_{j+1}(s,a) &> [\![\max(V_{j+1}[s'], R(s,a)) - K]\!] \\ &= [\![\max(V_j[s], R(s,a)) - K]\!]. \end{aligned}$$

  Combined with $V_j(s,a) > V_{j+1}(s,a)$ (see above), we obtain

  $$V_j(s,a) > [\![\max(V_j[s], R(s,a)) - K]\!],$$

  and, recalling the assumption $s' = s$, therefore $(s,a) \in \mathit{viol}(V_j)$.

- Suppose $a \notin \mathit{pref}(s, V_j)$ and $s' \neq s$. The latter implies $V_{j+1}[s'] = V_j[s']$. Since $(s,a) \in \mathit{viol}(V_{j+1})$, we have

  $$\begin{aligned} V_{j+1}(s,a) &> [\![\max(V_{j+1}[s'], R(s,a)) - K]\!] \\ &= [\![\max(V_j[s'], R(s,a)) - K]\!]. \end{aligned}$$

  Combined with $V_j(s,a) > V_{j+1}(s,a)$ (see above), we obtain

  $$V_j(s,a) > [\![\max(V_j[s'], R(s,a)) - K]\!],$$

  and therefore $(s,a) \in \mathit{viol}(V_j)$.

Second case Suppose $(s,a) \notin \mathit{viol}(V_{j+1})$. Hence,

$$V_{j+1}(s,a) \leq [\![\max(V_{j+1}[s'], R(s,a)) - K]\!] .$$

In order for $(s,a) \in \mathit{viol}(V^*)$, which we assumed to be true, it is necessary that the value of s' is strictly decreased before the end of fragment F. Otherwise, for all configuration indices $k \geq j+1$ in fragment F, we would have $V_{j+1}[s'] \leq V_k[s']$; and, combined with the assumption that transition $(s_j, V_j) \xrightarrow{a_j,\, s_{j+1}} (s_{j+1}, V_{j+1})$ is the last transition of F in which (s,a) is updated, we obtain

$$\begin{aligned} V_k(s,a) &= V_{j+1}(s,a) \\ &\leq [\![\max(V_{j+1}[s'], R(s,a)) - K]\!] \\ &\leq [\![\max(V_k[s'], R(s,a)) - K]\!] , \end{aligned}$$

implying $(s,a) \notin \mathit{viol}(V^*)$, which is false.

So, still inside fragment F, we can consider the *first* transition after configuration $j+1$ where the value of s' is strictly decreased:

$$(s_k, V_k) \xrightarrow{a_k,\, s_{k+1}} (s_{k+1}, V_{k+1}),$$

where $k \geq j+1$ and $V_{k+1}[s'] < V_k[s']$. This implies $s_k = s'$.[13] Denote $\delta(s', a_k) = \{s''\}$. Now, by Algorithm 1,

$$V_{k+1}(s', a_k) = [\![V_k(s', a_k) + \max(V_k[s''], R(s', a_k)) - K - V_k[s']]\!] .$$

Also, we have $a_k \in \mathit{pref}(s', V_k)$; otherwise there would be some action $a' \in \mathit{pref}(s', V_k)$ with $a' \neq a_k$ and $V_k[s'] = V_k(s', a') = V_{k+1}(s', a')$, implying $V_k[s'] \leq V_{k+1}[s']$, which is false. Since $a_k \in \mathit{pref}(s', V_k)$, we have $V_k[s'] = V_k(s', a_k)$. This can be used to simplify the above equation for $V_{k+1}(s', a_k)$, as follows:

$$V_{k+1}(s', a_k) = [\![\max(V_k[s''], R(s', a_k)) - K]\!] .$$

Next, since $V_k[s'] > V_{k+1}[s']$, we observe

$$V_k(s', a_k) = V_k[s'] > V_{k+1}[s'] \geq V_{k+1}(s', a_k).$$

In combination with the simplified equation for $V_{k+1}(s', a_k)$, we obtain

$$V_k(s', a_k) > [\![\max(V_k[s''], R(s', a_k)) - K]\!] .$$

Therefore, $(s', a_k) \in \mathit{viol}(V_k)$.

Now, Property B.3 (below) gives us $W < V_{j+1}[s']$. Since $(s_k, V_k) \xrightarrow{a_k,\, s_{k+1}} (s_{k+1}, V_{k+1})$ is the first transition after configuration $j+1$ with a value decrement on state s', we have $V_{j+1}[s'] \leq V_k[s']$. Combined, $W < V_k[s']$. Since $V_k[s'] = V_k(s', a_k)$ (see above), we obtain $W < V_k(s', a_k)$.

Overall, we obtain $\mathit{viol\text{-}max}(V_k) > W = \mathit{viol\text{-}max}(V^*) = \mathit{viol\text{-}max}(V_i)$, which contradicts Property B.1 (since $i \leq k$). This is the desired contradiction.

Property B.3. We have $W < V_{j+1}[s']$.

[13] If $s_k \neq s'$ then always $V_{k+1}[s'] = V_k[s']$.

Proof. Below we show that $V_{j+1}(s,a) \leq V_{j+1}[s'] - K$. Therefore,

$$W = V_{j+1}(s,a) < V_{j+1}(s,a) + K \leq V_{j+1}[s'],$$

giving $W < V_{j+1}[s']$, as desired.

We are left to show $V_{j+1}(s,a) \leq V_{j+1}[s'] - K$. First, it must be $V_{j+1}(s,a) > [\![R(s,a) - K]\!]$. Otherwise, considering $V_{j+1}(s,a) \leq [\![R(s,a) - K]\!]$, since the transition $(s_j, V_j) \xrightarrow{a_j,\, s_{j+1}} (s_{j+1}, V_{j+1})$ is the last transition in the run-fragment F where (s,a) is executed, for all configuration indices $k \geq j+1$ in F, we would have

$$\begin{aligned} V_k(s,a) &= V_{j+1}(s,a) \\ &\leq [\![R(s,a) - K]\!] \\ &\leq \max([\![V_k[s'] - K]\!], [\![R(s,a) - K]\!]) \\ &= [\![\max(V_k[s'], R(s,a)) - K]\!]. \end{aligned}$$

In particular, $V^*(s,a) \leq [\![\max(V^*[s'], R(s,a)) - K]\!]$, implying $(s,a) \notin \mathit{viol}(V^*)$, which is false.

Subsequently, $V_{j+1}(s,a) > [\![R(s,a) - K]\!]$ implies $V_{j+1}(s,a) \leq [\![V_{j+1}[s'] - K]\!]$. Otherwise,

$$\begin{aligned} V_{j+1}(s,a) &> \max([\![V_{j+1}[s'] - K]\!], [\![R(s,a) - K]\!]) \\ &= [\![\max(V_{j+1}[s'], R(s,a)) - K]\!], \end{aligned}$$

implying $(s,a) \in \mathit{viol}(V_{j+1})$, which we assumed to be false.

Lastly, since $V_{j+1}(s,a) = W > 0$, we know $[\![V_{j+1}[s'] - K]\!] > 0$. Therefore, $V_{j+1}[s'] - K > 0$. Hence, we may write $V_{j+1}(s,a) \leq V_{j+1}[s'] - K$. □

B.2 Proof of Property 3.17

Let V be a valid value function. Towards a contradiction, suppose there is some $s \in S$ with $V[s] > \mathit{opt\text{-}val}(s)$. By Property B.4 (below), there is an action $a \in \mathit{pref}(s,V)$, denoting $\delta(s,a) = \{s'\}$, with $V[s'] > \mathit{opt\text{-}val}(s')$ and $V[s] < V[s']$. Property B.4 can subsequently be applied to s'. By repeatedly applying Property B.4, we can establish an infinite sequence of the following form:

$$s_1 \xrightarrow{a_1} s_2 \xrightarrow{a_2} \dots,$$

where $V[s_i] < V[s_{i+1}]$ for each $i \geq 1$. But since there are a finite number of states, there must be two indices j and k with $j < k$ and $s_j = s_k$. Then $V[s_j] < V[s_k]$ is the desired contradiction.

Property B.4. Let V be the considered valid value function. Let $s \in S$. If $V[s] > \mathit{opt\text{-}val}(s)$ then $\exists a \in \mathit{pref}(s,V)$, denoting $\delta(s,a) = \{s'\}$, with

- $V[s'] > \mathit{opt\text{-}val}(s')$; and,
- $V[s] < V[s']$.

Proof. Let $a \in \mathit{pref}(s,V)$ be arbitrary, and denote $\delta(s,a) = \{s'\}$.

First we show that $V[s'] > \textit{opt-val}(s')$. Always,[14]

$$\begin{aligned}\textit{opt-val}(s) &\geq \max(\llbracket \textit{opt-val}(s') - K \rrbracket, \llbracket R(s,a) - K \rrbracket) && \text{(B.2)}\\ &= \llbracket \max(\textit{opt-val}(s'), R(s,a)) - K \rrbracket . && \text{(B.3)}\end{aligned}$$

Towards a contradiction, suppose $V[s'] \leq \textit{opt-val}(s')$. Then, using all assumptions (including validity of V), and the equality $V[s] = V(s,a)$ (by $a \in \textit{pref}(s,V)$), we have

$$\begin{aligned}\textit{opt-val}(s) &< V[s]\\ &= V(s,a)\\ &\leq \llbracket \max(V[s'], R(s,a)) - K \rrbracket\\ &\leq \llbracket \max(\textit{opt-val}(s'), R(s,a)) - K \rrbracket ,\end{aligned}$$

which contradicts Equation (B.3). Therefore $V[s'] > \textit{opt-val}(s')$.

Now we show that $V[s] < V[s']$. Since $V[s] > \textit{opt-val}(s)$ by assumption, Equation (B.2) implies $V[s] > \llbracket R(s,a) - K \rrbracket$. Together with $V[s] = V(s,a)$ (by $a \in \textit{pref}(s,V)$) and validity, we have

$$\begin{aligned}\llbracket R(s,a) - K \rrbracket &< V[s]\\ &= V(s,a)\\ &\leq \llbracket \max(V[s'], R(s,a)) - K \rrbracket\\ &= \max(\llbracket V[s'] - K \rrbracket, \llbracket R(s,a) - K \rrbracket). && \text{(B.4)}\end{aligned}$$

We have $\llbracket V[s'] - K \rrbracket > \llbracket R(s,a) - K \rrbracket$ because otherwise Equation (B.4) would imply the contradiction $\llbracket R(s,a) - K \rrbracket < \llbracket R(s,a) - K \rrbracket$. Thus $\llbracket V[s'] - K \rrbracket > 0$, causing $V[s'] - K > 0$, and therefore $\llbracket V[s'] - K \rrbracket = V[s'] - K$. Validity now implies,

$$\begin{aligned}V[s] &= V(s,a)\\ &\leq \max(\llbracket V[s'] - K \rrbracket, \llbracket R(s,a) - K \rrbracket)\\ &= \llbracket V[s'] - K \rrbracket\\ &= V[s'] - K.\end{aligned}$$

Hence, $V[s] \leq V[s'] - K$ and therefore $V[s] < V[s']$ (using that $K \geq 1$): $V[s] < V[s] + K \leq V[s']$. □

B.3 Proof of Property 3.19

If $|p| = 1$ then necessarily $\textit{opt-val}(s) = \textit{val}(p) = \llbracket R(s,a) - K \rrbracket$.

Henceforth, we assume $|p| \geq 2$. Let p' denote the suffix of p after omitting the first pair (s,a). By Lemma 3.22,

$$\textit{val}(p) = \max(\llbracket R(s,a) - K \rrbracket, \llbracket \textit{val}(p') - K \rrbracket).$$

[14] Note that $\textit{opt-val}(s) \geq \llbracket R(s,a) - K \rrbracket$ because (s,a) is an action-path for s. Also, $\textit{opt-val}(s) \geq \llbracket \textit{opt-val}(s') - K \rrbracket$ because any optimal action-path for s' can be extended to an action-path for s by adding (s,a) to the front. Formally, letting p' be an action-path for s' with $\textit{val}(p') = \textit{opt-val}(s')$, and letting p be the extension of p' by adding (s,a) to the front; Lemma 3.22 implies $\textit{opt-val}(s) \geq \textit{val}(p) = \max(\llbracket R(s,a) - K \rrbracket, \llbracket \textit{opt-val}(s') - K \rrbracket)$.

If $[\![R(s,a)-K]\!] \geq [\![val(p')-K]\!]$ then again $opt\text{-}val(s) = [\![R(s,a)-K]\!]$. Henceforth we assume $[\![R(s,a)-K]\!] < [\![val(p')-K]\!]$. This implies $[\![val(p')-K]\!] > 0$, causing $val(p')-K > 0$, so we write more simply

$$val(p) = val(p') - K.$$

We now show concretely that $val(p') = opt\text{-}val(s')$, giving, as desired

$$opt\text{-}val(s) = opt\text{-}val(s') - K.$$

We separately show $val(p') \leq opt\text{-}val(s')$ and $opt\text{-}val(s') \leq val(p')$.

Direction 1. Since p' is an action-path for s' we observe

$$val(p') \leq opt\text{-}val(s').$$

Direction 2. Next, let p'' be an action-path for s' with $val(p'') = opt\text{-}val(s')$. We can add the pair (s,a) to the front of p'', resulting in a path p'''. By Lemma 3.22,

$$\begin{aligned} val(p''') &= \max([\![R(s,a)-K]\!], [\![val(p'')-K]\!]) \\ &\geq [\![val(p'')-K]\!]. \end{aligned}$$

Also, by definition of $opt\text{-}val(s)$, we have $val(p''') \leq opt\text{-}val(s) = val(p) = val(p') - K$. Everything combined, we have

$$[\![val(p'')-K]\!] \leq val(p') - K.$$

We have $val(p'') \leq val(p')$: otherwise, considering $val(p'') > val(p')$, since $val(p') - K > 0$ (see above), we would have $val(p'') - K > 0$; and subsequently $[\![val(p'')-K]\!] = val(p'') - K > val(p') - K$, which is false.

Now, $val(p'') \leq val(p')$, combined with $val(p'') = opt\text{-}val(s')$, implies the second direction that was sought:

$$opt\text{-}val(s') \leq val(p').$$

C Proof of Lemma 3.22

By definition of path-value (Equation (3.1)),

$$val(p) = \max\{[\![R(s_1,a_1)-K]\!], [\![R(s_2,a_2)-2K]\!], \ldots, [\![R(s_n,a_n)-nK]\!]\}.$$

We may rewrite this as follows:

$$val(p) = \max([\![R(s_1,a_1)-K]\!], m),$$

where

$$\begin{aligned} m &= \max\{[\![R(s_2,a_2)-2K]\!], \ldots, [\![R(s_n,a_n)-nK]\!]\} \\ &= [\![\max\{R(s_2,a_2)-2K, \ldots, R(s_n,a_n)-nK\}]\!]. \end{aligned}$$

Subsequently,

$$\begin{aligned} m &= [\![\max\{R(s_2,a_2)-K, \ldots, R(s_n,a_n)-(n-1)K\} - K]\!]. \\ &= [\![[\![\max\{R(s_2,a_2)-K, \ldots, R(s_n,a_n)-(n-1)K\}]\!] - K]\!] \\ &= [\![\max\{[\![R(s_2,a_2)-K]\!], \ldots, [\![R(s_n,a_n)-(n-1)K]\!]\} - K]\!] \\ &= [\![val(p')-K]\!]. \end{aligned}$$

D Proof details of Theorem 4.6

D.1 Auxiliary general properties

Theorem 4.6 assumes that all initial values are below M. For a value function V, we define

$$\mathit{highest}(V) = \max\{V(s,a) \mid (s,a) \in S \times A\}.$$

The following property will be useful:

Property D.1. For any run on the task, for any encountered value function V, we have $\mathit{highest}(V) < M$.

Proof. We show the property by induction on the transitions of the run. By assumption, the property is true for the initial value function. Now, consider a transition

$$(s,V) \xrightarrow{a,\,s'} (s',V').$$

Assume $\mathit{highest}(V) < M$. We show $\mathit{highest}(V') < M$. For each $(s'',a'') \in S \times A$ with $(s'',a'') \neq (s,a)$ we have $V'(s'',a'') = V(s'',a'') \leq \mathit{highest}(V) < M$. For the pair (s,a) itself we have, by Algorithm 1,

$$\begin{aligned} V'(s,a) &= [\![V(s,a) + \big(\max(V[s'], R(s,a)) - K - V[s]\big)]\!] \\ &\leq [\![V[s] + \big(\max(V[s'], R(s,a)) - K - V[s]\big)]\!] \\ &= [\![\max(V[s'], R(s,a)) - K]\!]. \end{aligned}$$

By subsequently using $V[s'] \leq \mathit{highest}(V) < M$ and $R(s,a) \leq M$, we have

$$V'(s,a) \leq [\![M - K]\!].$$

Lastly, we use $M - K > 0$ by Property D.2 (below), which implies $[\![M-K]\!] = M - K$, to obtain

$$V'(s,a) \leq M - K < M.$$

□

Property D.2. We have $M - K > 0$.

Proof. Since $|S| \geq 1$ we have $M - K \geq M - |S|\,K$. And $M - |S|\,K > 0$ because the task is a navigation problem. □

D.2 Auxiliary properties of strategies

Let V be a value function. For uniformity, we define $z_0(V) = \emptyset$. We define $\mathit{fixp}(V)$ as the smallest index $n \in \mathbb{N}$ for which $z_n(V) = z_k(V)$ for all $k \geq n$, i.e., $\mathit{fixp}(V)$ is the fixpoint index. Possibly $\mathit{fixp}(V) = 0$, when no states can be added to the strategy.

Property D.3. Let V be a value function. For each $s \in \mathit{strategy}(V)$, we have $V[s] \geq M - \mathit{fixp}(V)K$.

Proof. Denote $n = \mathit{fixp}(V)$. Let $s \in \mathit{strategy}(V)$. There is a smallest index $j \geq 1$ with $s \in z_j(V)$, implying $V[s] = M - jK$. We have $j \leq n$: otherwise, considering $j > n$, we would have $s \in z_j(V) \setminus z_n(V)$, which is not possible because $n = \mathit{fixp}(V)$.

Now, $j \leq n$ implies $V[s] \geq M - nK$. □

Property D.4. Let V be a value function. We have

$$\mathit{fixp}(V) \leq |\mathit{strategy}(V)|\,.$$

Proof. Abbreviate $n = \mathit{fixp}(V)$. If $n = 0$ then the property is immediately true. Henceforth, suppose $n \geq 1$. We show by induction on $j = n, \ldots, 1$ that

$$\exists s \in z_j(V) \text{ with } V[s] = M - jK.$$

For any two states s and s', if $V[s] \neq V[s']$ then $s \neq s'$; hence the inductive property implies $n \leq |\mathit{strategy}(V)|$, as desired.

- Base case: $j = n$. By choice of n as the smallest index after which no more states are added to the strategy, we have $z_n(V) \neq z_{n-1}(V)$. Hence, $z_n(V)$ extends $z_{n-1}(V)$ with at least one state s satisfying $V[s] = M - nK$.
- Inductive step. Let $j \geq 2$, with the assumption that $z_j(V)$ contains a state s with $V[s] = M - jK$. By definition of $z_j(V)$, for each $a \in \mathit{pref}(s, V)$ there must be some state $s' \in \delta(s, a) \subseteq z_{j-1}(V)$ with $V[s'] = M-(j-1)K$. Hence, there is at least one state $s' \in z_{j-1}(V)$ with $V[s'] = M - (j-1)K$.

□

Property D.5. Let V be a value function. If $\mathit{strategy}(V) \neq S$ then for each $s \in \mathit{strategy}(V)$ we have $V[s] > K$.

Proof. Denote $n = \mathit{fixp}(V)$. Since $\mathit{strategy}(V) \neq S$, and yet always $\mathit{strategy}(V) \subseteq S$, we have $|\mathit{strategy}(V)| < |S|$. Combined with $n \leq |\mathit{strategy}(V)|$ (by Property D.4), we see that

$$n + 1 \leq |S|\,.$$

Since $M - |S|\,K > 0$ by assumption on navigation problems, we obtain

$$M - (n+1)K \geq M - |S|\,K > 0,$$

resulting in $M - nK > K$.

Now, let $s \in \mathit{strategy}(V)$. Because $V[s] \geq M - nK$ by Property D.3, we now observe, as desired,

$$V[s] > K.$$

□

Property D.6. Let V be a value function. For each $s \in \mathit{strategy}(V)$ we have

$$0 < V[s] \leq M - K.$$

Proof. For the upper bound, we note that for each $s \in z_1(V)$ we have $V[s] = M - K$, and for each $i \geq 2$, for each $s \in z_i(V) \setminus z_{i-1}(V)$, we have $V[s] = M - iK < M - K$.

For the lower bound, let $s \in \mathit{strategy}(V)$. We first recall that $\mathit{fixp}(V) \leq |\mathit{strategy}(V)|$ by Property D.4. Combined with $|\mathit{strategy}(V)| \leq |S|$ (which is always true), we arrive at $\mathit{fixp}(V) \leq |S|$. Now, combined with Property D.3, and the assumption $M - |S|\,K > 0$ on navigation problems, we obtain, as desired,

$$V[s] \geq M - \mathit{fixp}(V)K \geq M - |S|\,K > 0.$$

□

Property D.7. Let V be a value function. We have

$$\mathit{strategy}(V) \subseteq \mathit{reduce}(T).$$

Proof. We show by induction on $j = 1, 2, \ldots$ that $z_j(V) \subseteq \mathit{reduce}(T)$.

- For the base case, we know for each $s \in z_1(V)$ that $\exists a \in \mathit{pref}(s, V)$ with $(s, a) \in \mathit{rewards}(T)$. Therefore $z_1(V) \subseteq \mathit{goals}(T) \subseteq \mathit{reduce}(T)$.
- For the inductive step, let $j \geq 1$ and assume $z_j(V) \subseteq \mathit{reduce}(T)$. We show that $z_{j+1}(V) \subseteq \mathit{reduce}(T)$. Suppose $z_j(V) \subsetneq z_{j+1}(V)$. Let $s \in z_{j+1}(V) \setminus z_j(V)$. By definition of $z_{j+1}(V)$ we know $\exists a \in \mathit{pref}(s, V)$ with $\delta(s, a) \subseteq z_j(V)$. By applying the induction hypothesis, we know $\delta(s, a) \subseteq \mathit{reduce}(T)$. Denoting $\mathit{reduce}(T) = \bigcup_{i=1}^{\infty} L_i(T)$, we can consider an index k with $\delta(s, a) \subseteq L_k(T)$. Hence $s \in L_{k+1}(T) \subseteq \mathit{reduce}(T)$.

□

D.3 Proof of Property 4.10

Consider the suffix $\mathcal{X}$ of a greedy run,

$$(s_1, V_1) \xrightarrow{a_1, s_2} (s_2, V_2) \xrightarrow{a_2, s_3} \ldots$$

where (s_1, V_1) is a good configuration. In the suffix, consider a finite sequence of transitions forming a state cycle, denoted as

$$(s_j, V_j) \xrightarrow{a_j,\, s_{j+1}} \ldots \xrightarrow{a_{j+(n-1)},\, s_{j+n}} (s_{j+n}, V_{j+n}),$$

where $n \geq 1$ and $s_{j+n} = s_j$. Towards a contradiction, if none of the transitions between (s_j, V_j) and (s_{j+n}, V_{j+n}) has reward then Property D.8 (below) tells us that, inside value function V_1,

$$V_1[s_j] < \ldots < V_1[s_{j+n}] = V_1[s_j],$$

which is a contradiction. Hence, all state cycles in $\mathcal{X}$ contain reward.

Property D.8. Consider the suffix of a greedy run,

$$(s_1, V_1) \xrightarrow{a_1, s_2} (s_2, V_2) \xrightarrow{a_2, s_3} \ldots$$

where (s_1, V_1) is a good configuration. For each $i \geq 1$, if $(s_i, a_i) \notin \mathit{rewards}(T)$ then $V_1[s_i] < V_1[s_{i+1}]$. Note the special role played by V_1. *(Proof in Appendix D.3.1.)* □

D.3.1 Proof of Property D.8

Regarding notation, for any two value functions V and V', we write $V \triangleright V'$ if for each $s \in \mathit{strategy}(V)$ the following conditions are satisfied:

1. $\forall a \in \mathit{pref}(s, V)$: $V[s] \leq V'(s, a) \leq M - K$; and,
2. $\forall a \in A \setminus \mathit{pref}(s, V)$: $V'(s, a) < V[s]$.

We note that always $V \triangleright V$.[15] One may read $V \triangleright V'$ as V *causes* V', because for the states in *strategy*(V) the action-preference in V' is strongly related to the action-preference in V; see also Property D.9 below.

Property D.9. Let V and V' be two value functions with $V \triangleright V'$. For each $s \in \mathit{strategy}(V)$ we have $\mathit{pref}(s, V') \subseteq \mathit{pref}(s, V)$.

Proof. Let $s \in \mathit{strategy}(V)$. Abbreviate $N = A \setminus \mathit{pref}(s, V)$. Below, we show $N \cap \mathit{pref}(s, V') = \emptyset$. Then,

$$\begin{aligned} \mathit{pref}(s, V') &= \mathit{pref}(s, V') \setminus N \\ &\subseteq A \setminus N \\ &= \mathit{pref}(s, V). \end{aligned}$$

If $N = \emptyset$ then immediately $N \cap \mathit{pref}(s, V') = \emptyset$. Henceforth, suppose $N \neq \emptyset$. Let $a_1 \in N$ and $a_2 \in \mathit{pref}(s, V)$.[16] Because $V \triangleright V'$, we have

$$V'(s, a_1) < V[s] \leq V'(s, a_2) \leq V'[s].$$

Hence, $V'(s, a_1) < V'[s]$, giving $a_1 \notin \mathit{pref}(s, V')$. □

We continue with the proof of Property D.8. Abbreviate $V = V_1$. Consider a transition

$$(s_i, V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1}, V_{i+1}),$$

with $(s_i, a_i) \notin \mathit{rewards}(T)$. By Property D.10 (below), we know $s_i \in \mathit{strategy}(V)$ and $V \triangleright V_i$. Subsequently, Property D.9 gives $\mathit{pref}(s_i, V_i) \subseteq \mathit{pref}(s_i, V)$. Since $a_i \in \mathit{pref}(s_i, V_i)$ by greediness of the run, we find $a_i \in \mathit{pref}(s_i, V)$.

Next, because $s_i \in \mathit{strategy}(V)$, we can consider the smallest index j satisfying $s_i \in z_j(V)$, which implies $V[s_i] = M - jK$. Since $a_i \in \mathit{pref}(s_i, V)$ (see above) and $(s_i, a_i) \notin \mathit{rewards}(T)$, we have $j \geq 2$. Then, by definition of $z_j(V)$, we have $\delta(s_i, a_i) \subseteq z_{j-1}(V)$. In particular, since $s_{i+1} \in \delta(s_i, a_i)$, we see $s_{i+1} \in z_{j-1}(V)$. Therefore $V[s_{i+1}] \geq M - (j-1)K > M - jK = V[s_i]$. Overall, $V[s_i] < V[s_{i+1}]$, as desired.[17]

Property D.10. Consider the suffix of a greedy run,

$$(s_1, V_1) \xrightarrow{a_1, s_2} (s_2, V_2) \xrightarrow{a_2, s_3} \dots$$

where (s_1, V_1) is a good configuration. Abbreviating, $V = V_1$, for each $i \geq 1$ we have

1. $s_i \in \mathit{strategy}(V)$;
2. $V \triangleright V_i$.

(Proof in Appendix D.3.2.) □

[15] Let $s \in \mathit{strategy}(V)$. For any $a \in \mathit{pref}(s, V)$, always $V[s] = V(s, a)$; and, $V[s] \leq M - K$ by Property D.6. For any $a \in A \setminus \mathit{pref}(s, V)$, always $V(s, a) < V[s]$.

[16] Note that always $\mathit{pref}(s, V) \neq \emptyset$.

[17] For completeness, we note that not necessarily $V_i[s_i] < V_i[s_{i+1}]$.

D.3.2 Proof of Property D.10

We show these properties by induction on $i = 1, 2, \ldots$. For the base case, $i = 1$, we note the following:

1. We have $s_1 \in \mathit{strategy}(V_1) = \mathit{strategy}(V)$ since (s_1, V_1) is a good configuration.
2. Always $V_1 \triangleright V_1$.

For the inductive step, consider a transition,

$$(s_i, V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1}, V_{i+1}),$$

where $i \geq 1$. As induction hypothesis, we assume $s_i \in \mathit{strategy}(V)$ and $V \triangleright V_i$. We show that the induction properties are also true for (s_{i+1}, V_{i+1}).

First property We show that $s_{i+1} \in \mathit{strategy}(V)$. If $(s_i, a_i) \in \mathit{rewards}(T)$ then $s_{i+1} \in S^{\mathrm{start}}$ since task T is restartable. Moreover, because $S^{\mathrm{start}} \subseteq \mathit{strategy}(V)$ by goodness of (s_1, V_1), we obtain $s_{i+1} \in \mathit{strategy}(V_1) = \mathit{strategy}(V)$.

Suppose $(s_i, a_i) \notin \mathit{rewards}(T)$. Since $s_i \in \mathit{strategy}(V)$ by the induction hypothesis, we can consider the smallest index j with $s_i \in z_j(V)$. Also, $V \triangleright V_i$ by the induction hypothesis. Subsequently, $\mathit{pref}(s_i, V_i) \subseteq \mathit{pref}(s_i, V)$ by Property D.9. Since $a_i \in \mathit{pref}(s_i, V_i)$ by greediness of the run, we find $a_i \in \mathit{pref}(s_i, V)$. The assumption $(s_i, a_i) \notin \mathit{rewards}(T)$ now implies $j \geq 2$. By definition of $z_j(V)$ with $j \geq 2$, we know that $\delta(s_i, a_i) \subseteq z_{j-1}(V) \subseteq \mathit{strategy}(V)$. In particular, $s_{i+1} \in \mathit{strategy}(V)$.

Second property We show that $V \triangleright V_{i+1}$. Let $s \in \mathit{strategy}(V)$. As above, let j be the smallest index for which $s \in z_j(V)$.

Preferred actions. Let $a \in \mathit{pref}(s, V)$. We have to show that

$$V[s] \leq V_{i+1}(s, a) \leq M - K.$$

If $(s, a) \neq (s_i, a_i)$ then $V_{i+1}(s, a) = V_i(s, a)$, and the induction hypothesis $V \triangleright V_i$ implies $V[s] \leq V_{i+1}(s, a) \leq M - K$. Henceforth, suppose $(s, a) = (s_i, a_i)$. By Algorithm 1,

$$V_{i+1}(s, a) = [\![V_i(s, a) + (\max(V_i[s_{i+1}], R(s, a)) - K - V_i[s])]\!].$$

Since $a = a_i \in \mathit{pref}(s, V_i)$ by greedy action selection, we have $V_i[s] = V_i(s, a)$, and the expression simplifies to

$$V_{i+1}(s, a) = [\![\max(V_i[s_{i+1}], R(s, a)) - K]\!].$$

Another general observation, is that $V_i[s_{i+1}] \leq M - K$: since $s_{i+1} \in \mathit{strategy}(V)$ (see above), and $V \triangleright V_i$ (by the induction hypothesis), we have

- for each $a' \in \mathit{pref}(s_{i+1}, V)$: $V[s_{i+1}] \leq V_i(s_{i+1}, a') \leq M - K$;
- for each $a' \in A \setminus \mathit{pref}(s_{i+1}, V)$: $V_i(s_{i+1}, a') < V[s_{i+1}] \leq M - K$ (using Property D.6 for the upper bound).

Next, we distinguish between two cases, as follows.

- Suppose $(s,a) \in \mathit{rewards}(T)$. From the proof of the first induction property above, we recall that $a \in \mathit{pref}(s,V)$. Now, $(s,a) \in \mathit{rewards}(T)$ implies $s \in z_1(V)$. Therefore $V[s] = M - K$. Also, recall that $M - K > 0$ (Property D.2).

  By applying $V_i[s_{i+1}] \leq M-K$ (see above) and $R(s,a) = M$ (since $(s,a) \in \mathit{rewards}(T)$) to the equation for $V_{i+1}(s,a)$, we obtain:[18]

$$\begin{aligned} V_{i+1}(s,a) &= [\![\max(V_i[s_{i+1}], R(s,a)) - K]\!] \\ &= [\![M - K]\!] \\ &= M - K \\ &= V[s]. \end{aligned}$$

  Hence, $V[s] \leq V_{i+1}(s,a) \leq M - K$.

- Suppose $(s,a) \notin \mathit{rewards}(T)$. From the proof of the first induction property above, we recall that $a \in \mathit{pref}(s,V)$. Now, $(s,a) \notin \mathit{rewards}(T)$ implies $s \in z_j(V)$ with $j \geq 2$. By applying $R(s,a) = 0$, and using $V_i[s_{i+1}] \geq 0$, the earlier equation of $V_{i+1}(s,a)$ is simplified as follows:

$$V_{i+1}(s,a) = [\![V_i[s_{i+1}] - K]\!]\,.$$

  Before we continue, we show $V[s_{i+1}] \leq V_i[s_{i+1}]$. Since $a \in \mathit{pref}(s,V)$, the definition of $s \in z_j(V)$ with $j \geq 2$ implies $\delta(s,a) \subseteq z_{j-1}(V)$. In particular, $s_{i+1} \in z_{j-1}(V) \subseteq \mathit{strategy}(V)$, which implies $V[s_{i+1}] \geq M - (j-1)K$. Moreover, combining $s_{i+1} \in \mathit{strategy}(V)$ and $V \triangleright V_i$, and letting $a' \in \mathit{pref}(s_{i+1}, V)$, we have

$$V[s_{i+1}] \leq V_i(s_{i+1}, a') \leq V_i[s_{i+1}].$$

  Next, $s \in z_j(V)$ gives $V[s] = M - jK$, and $j \geq 2$ further implies,

$$\begin{aligned} V[s] &= M - (j-1)K - K \\ &\leq V[s_{i+1}] - K \\ &\leq V_i[s_{i+1}] - K. \end{aligned}$$

  Lastly, by applying the deduced inequalities $V[s] \leq V_i[s_{i+1}] - K$ and $V_i[s_{i+1}] \leq M-K$ (see earlier) to the last simplified equation for $V_{i+1}(s,a)$, and using $V[s] \geq 0$, we obtain:

$$V_{i+1}(s,a) = [\![V_i[s_{i+1}] - K]\!] \geq [\![V[s]]\!] = V[s],$$

  and, using $M - K > 0$,

$$V_{i+1}(s,a) = [\![V_i[s_{i+1}] - K]\!] \leq [\![M - K - K]\!] \leq [\![M - K]\!] = M - K.$$

Non-preferred actions. Let $a \in A \setminus \mathit{pref}(s,V)$. We have to show that

$$V_{i+1}(s,a) < V[s].$$

Recall $V \triangleright V_i$ by the induction hypothesis. We distinguish between two cases, as follows:

[18] Note in particular that $\max(M-K, M) = M$.

- Suppose $s \neq s_i$. We have $V_{i+1}(s,a) = V_i(s,a) < V[s]$, where the inequality is given by $V \triangleright V_i$.
- Suppose $s = s_i$. Since $s_i \in \mathit{strategy}(V)$, $V \triangleright V_i$, and $a_i \in \mathit{pref}(s_i, V_i)$ (by greediness of the run), Property D.9 tells us $a_i \in \mathit{pref}(s_i, V) = \mathit{pref}(s, V)$. Therefore, $a \neq a_i$, and again $V_{i+1}(s,a) = V_i(s,a) < V[s]$ (with the same reasoning as in the previous case).

D.4 Auxiliary properties of function β

Property D.11. Consider a transition generated by function β,

$$(s_i, V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1}, V_{i+1}).$$

For each $(s,a) \in \mathit{strategy}(V_i) \times A$, we have $V_{i+1}(s,a) = V_i(s,a)$. In words: no changes occur to the value of state-action pairs where the state is in the strategy.

Proof. Let $(s,a) \in \mathit{strategy}(V_i) \times A$. If $(s,a) \neq (s_i, a_i)$ then the value could not have changed during the transition.

Henceforth, suppose $(s,a) = (s_i, a_i)$. We show that $V_{i+1}(s_i, a_i) = V_i(s_i, a_i)$. Since $s_i \in \mathit{strategy}(V_i)$, we can consider the smallest index j with $s_i \in z_j(V_i)$. This implies $V_i[s_i] = M - jK$. Also, we have $a_i \in \mathit{pref}(s_i, V_i)$ because function β only selects an action that the agent prefers. We distinguish between the following cases:

- Suppose $(s_i, a_i) \in \mathit{rewards}(T)$. Therefore $j = 1$, and subsequently $V_i[s_i] = M - K$. Next, by Algorithm 1, we have

  $$V_{i+1}(s_i, a_i) = [\![V_i(s_i, a_i) + (\max(V_i[s_{i+1}], R(s_i, a_i)) - K - V_i[s_i])]\!]\,.$$

  We have $V_i[s_{i+1}] < M$ by Property D.1, and $R(s_i, a_i) = M$ since $(s_i, a_i) \in \mathit{rewards}(T)$. Overall,

  $$\begin{aligned} V_{i+1}(s_i, a_i) &= [\![V_i(s_i, a_i) + (M - K) - (M - K)]\!] \\ &= [\![V_i(s_i, a_i)]\!] \\ &= V_i(s_i, a_i), \end{aligned}$$

  where the last step uses $V_i(s_i, a_i) \geq 0$.
- Suppose $(s_i, a_i) \notin \mathit{rewards}(T)$. This implies $j \geq 2$. By definition of $z_j(V_i)$, we know
  1. $\delta(s_i, a_i) \subseteq z_{j-1}(V_i)$, which gives $V_i[s'] \geq M - (j-1)K$ for each $s' \in \delta(s_i, a_i)$; and,
  2. $\exists s' \in \delta(s_i, a_i)$ with $V_i[s'] = M - (j-1)K$.

  Therefore, $\mathit{expect}(V_i, s_i, a_i) = M - (j-1)K$. By subsequently using that $V_i(s_i, a_i) = V_i[s_i]$ since $a_i \in \mathit{pref}(s_i, V_i)$, and using $V_i[s_i] = M - jK$, we see

  $$V_i(s_i, a_i) = \mathit{expect}(V_i, s_i, a_i) - K.$$

  We can now see that in the specification of function β, we have to exclude all cases except Case 2.2.3.2:

- Case 1 is not possible because $V_i[s_i] > 0$, as given by $s_i \in \mathit{strategy}(V_i)$ and Property D.6.
- Case 2.1 is not possible because $\delta(s_i, a_i) \subseteq z_{j-1}(V_i) \subseteq \mathit{strategy}(V_i)$.
- Case 2.2.1 is not possible because $(s_i, a_i) \notin \mathit{rewards}(T)$ by assumption.
- Case 2.2.2 is not possible because it would demand $V_i(s_i, a_i) \neq \mathit{expect}(V_i, s_i, a_i) - K$, which is false, as we have shown above.
- Case 2.2.3.1 is not possible because $(s_i, a_i) \notin \mathit{rewards}(T)$ by assumption.

Therefore, only Case 2.2.3.2 is possible. Importantly, Case 2.2.3.2 chooses $s_{i+1} \in \delta(s_i, a_i)$ to satisfy $V_i[s_{i+1}] = \mathit{expect}(V_i, s_i, a_i)$. We also have $R(s_i, a_i) = 0$. Lastly, since $a_i \in \mathit{pref}(s_i, V_i)$, we have $V_i[s_i] = V_i(s_i, a_i) = \mathit{expect}(V_i, s_i, a_i) - K$, as shown above. Subsequently, the equation for $V_{i+1}(s_i, a_i)$, as given by Algorithm 1, can be simplified in the following manner:

$$\begin{aligned} V_{i+1}(s_i, a_i) &= [\![V_i(s_i, a_i) + (\max(V_i[s_{i+1}], R(s_i, a_i)) - K - V_i[s_i])]\!] \\ &= [\![V_i(s_i, a_i) + (V_i[s_{i+1}] - K) - V_i[s_i])]\!] \\ &= [\![V_i(s_i, a_i) + (\mathit{expect}(V_i, s_i, a_i) - K) - (\mathit{expect}(V_i, s_i, a_i) - K)]\!] \\ &= [\![V_i(s_i, a_i)]\!] \\ &= V_i(s_i, a_i). \end{aligned}$$

where the last step uses $V_i(s_i, a_i) \geq 0$.

□

D.5 Proof of Property 4.11

Let s be the fixed start state. Consider a path

$$s_n \xrightarrow{a_n} s_{n-1} \ldots s_1 \xrightarrow{a_1} s_0,$$

where $n \geq 1$, $\{s_n, \ldots, s_1\} \subseteq \mathit{non\text{-}reduce}(T)$, and $s_0 = s$. We show by induction on $j = 1, \ldots, n$ that $s_j \in \mathit{dom}\left(g_j^{(s)}\right)$, which eventually implies $s_n \in \mathit{dom}\left(g_n^{(s)}\right) \subseteq \mathit{dom}\left(g^{(s)}\right)$. Because the task is reducible, every non-reducible state has a path inside $\mathit{non\text{-}reduce}(T)$ towards s. Hence, $\mathit{non\text{-}reduce}(T) \subseteq \mathit{dom}\left(g^{(s)}\right)$, as desired.

For the base case, we see $s_1 \in \mathit{dom}\left(g_1^{(s)}\right)$ because $s = s_0 \in \delta(s_1, a_1)$. For the inductive step, with $j \geq 2$ (and $j \leq n$), if not already $s_j \in \mathit{dom}\left(g_{j-1}^{(s)}\right)$ then surely $s_j \in \mathit{dom}\left(g_j^{(s)}\right)$ because $s_{j-1} \in \mathit{dom}\left(g_{j-1}^{(s)}\right)$ (by the induction hypothesis) and $s_{j-1} \in \delta(s_j, a_j)$.

D.6 Proof of Property 4.12

Let $\mathcal{X}$ denote the infinite sequence of transitions obtained by repeatedly applying function β starting at (s, V). By Property 4.13, under β, states are

never removed from the strategy, i.e., the strategy could in principle only grow. Towards a contradiction, suppose β is not able to eventually strictly extend the strategy, i.e., we have $\mathit{strategy}(V') = \mathit{strategy}(V)$ for all encountered value functions V' after (s, V).

By design, β jumps to a start state outside $\mathit{strategy}(V)$ during each reward transition (if possible). By Property 4.14, there are infinitely many reward transitions in $\mathcal{X}$, and since $S^{\text{start}} \not\subseteq \mathit{strategy}(V)$, we arrive infinitely often outside $\mathit{strategy}(V)$. There are two cases that could occur:

- There are infinitely many transitions where reward is obtained at a state outside $\mathit{strategy}(V)$.
- There are infinitely many transitions where reward is obtained at a state inside $\mathit{strategy}(V)$. This implies there are infinitely many transitions that jump from outside $\mathit{strategy}(V)$ to inside $\mathit{strategy}(V)$.

Because there are a finite number of configurations (Lemma 2.7), there are a finite number of possible transitions. Hence, in $\mathcal{X}$ there is either

- one particular transition, occurring infinitely often, where reward is obtained at a state outside $\mathit{strategy}(V)$; or,
- one particular transition, occurring infinitely often, that jumps from a state outside $\mathit{strategy}(V)$ to a state inside $\mathit{strategy}(V)$.

We distinguish between the two cases.

Reward outside *strategy*(V) Consider a transition specified by β,

$$(s_i, V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1}, V_{i+1}),$$

that occurs infinitely often in $\mathcal{X}$, and where $s_i \notin \mathit{strategy}(V)$ and $(s_i, a_i) \in \mathit{rewards}(T)$. We now analyze why function β has chosen (a_i, s_{i+1}), by looking at the specification of β.

- Case 1.1 is not possible: $s_i \in \mathit{non\text{-}reduce}(T)$ would contradict $(s_i, a_i) \in \mathit{rewards}(T)$.
- Case 1.2 is not possible: after the first execution of (s_i, a_i), the value of s_i will be at least $M - K$, which is strictly larger than zero (Property D.2); hence this case can not explain the infinite occurrences of the above transition.
- Case 1.3 is not possible: $s_i \notin \mathit{goals}(T)$ would contradict $(s_i, a_i) \in \mathit{rewards}(T)$.
- Case 2.1 is not possible; it would contradict $(s_i, a_i) \in \mathit{rewards}(T)$.
- Case 2.2.1 is not possible, as we now explain. We argue that after the first execution of (s_i, a_i), the value of (s_i, a_i) will remain $M - K$; hence Case 2.2.1 can not explain the infinite occurrences of the above transition. Consider a transition

$$(s_j, V_j) \xrightarrow{a_j,\, s_{j+1}} (s_{j+1}, V_{j+1}),$$

where $j \geq i$, and $(s_j, a_j) = (s_i, a_i)$. By Algorithm 1,

$$V_{j+1}(s_i, a_i) = [\![V_j(s_i, a_i) + \max(V_j[s_{j+1}], R(s_i, a_i)) - K - V_j[s_i]]\!] .$$

The equation can now be simplified as follows. Since $a_i = a_j \in \mathit{pref}(s_i, V_j)$ (since β always performs preferred actions), we have $V_j(s_i, a_i) = V_j[s_i]$. Moreover, $V_j[s_{j+1}] < M$ by Property D.1. Lastly, $R(s_i, a_i) = M$. We obtain the simplification,

$$V_{j+1}(s_i, a_i) = [\![M - K]\!] = M - K,$$

where we also use $M - K > 0$ (Property D.2).

- Case 2.2.2 is not possible; it would contradict $(s_i, a_i) \in \mathit{rewards}(T)$.
- Case 2.2.3.1 is possible. Recall that $s_i \notin \mathit{strategy}(V)$ by assumption. We now show $s_i \in \mathit{strategy}(V_i)$, implying $\mathit{strategy}(V) \subsetneq \mathit{strategy}(V_i)$; this is the desired contradiction.

  Concretely, we show $s_i \in z_1(V_i)$ (see Section 4.5.2).

  - We show $V_i[s_i] = M - K$. Because Case 2.2.1 was not applicable, $(s_i, a_i) \in \mathit{rewards}(T)$ and $a_i \in \mathit{pref}(s_i, V_i)$ (by design of β) together imply $V_i[s_i] = V_i(s_i, a_i) = M - K$.
  - Let $a \in \mathit{pref}(s_i, V_i)$. We show $(s_i, a) \in \mathit{rewards}(T)$. Towards a contradiction, if $(s_i, a) \notin \mathit{rewards}(T)$, since Case 2.2.2 was not applicable, we know $V_i(s_i, a) = \mathit{expect}(V_i, s_i, a) - K$. But $\mathit{expect}(V_i, s_i, a) < M$ by Property D.1, and thus $V_i(s_i, a) < M - K$, resulting in $a_i \notin \mathit{pref}(s_i, V_i)$, which is a contradiction.
- Case 2.2.3.2 is not possible; it would contradict $(s_i, a_i) \in \mathit{rewards}(T)$.

Arriving in *strategy*(V) Consider a transition

$$(s_i, V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1}, V_{i+1}),$$

that occurs infinitely often in $\mathcal{X}$, and where $s_i \notin \mathit{strategy}(V)$ and $s_{i+1} \in \mathit{strategy}(V)$. Note that $(s_i, a_i) \notin \mathit{rewards}(T)$: otherwise, we would jump (to a start state) outside $\mathit{strategy}(V)$.[19] We now analyze why function β has chosen (a_i, s_{i+1}), by looking at the specification of β.

- Case 1.1 is not possible. We would either jump to (1) a start state outside $\mathit{strategy}(V_i) = \mathit{strategy}(V)$ (which is possible because $S^{\text{start}} \not\subseteq \mathit{strategy}(V)$); or (2) a non-reducible state, which is also outside $\mathit{strategy}(V)$ by Property D.7. Either option would be impossible because $s_{i+1} \in \mathit{strategy}(V)$.
- Case 1.2 is not possible because $(s_i, a_i) \notin \mathit{rewards}(T)$.

[19] Also, the situation where $s_i \notin \mathit{strategy}(V)$ and $(s_i, a_i) \in \mathit{rewards}(T)$ was already discussed earlier.

- Case 1.3 is not possible, as we now explain. The case implies $V_i[s_i] = 0$. Therefore the value of s_i would have to be zero infinitely often in $\mathcal{X}$; we show this is not possible.

  We have $\delta(s_i, a_i) \subseteq \mathit{strategy}(V_i) = \mathit{strategy}(V)$: otherwise, Case 1.3 would have chosen a successor state s_{i+1} outside $\mathit{strategy}(V_i)$, which is false.

  Subsequently, noting $\mathit{strategy}(V) \neq S$ (as implied by $S^{\text{start}} \not\subseteq \mathit{strategy}(V)$), we apply Property D.5 to know $V[s'] > K$ for each $s' \in \delta(s_i, a_i)$. Hence, by Property D.12 (below), $V_i[s'] = V[s'] > K$ for each $s' \in \delta(s_i, a_i)$. During the above transition, from the viewpoint of Algorithm 1, we would have (using $s_{i+1} \in \delta(s_i, a_i)$ and $V_i[s_i] = 0$):

  $$\begin{aligned} d &= \max(V_i[s_{i+1}], R(s_i, a_i)) - K - V_i[s_i] \\ &= \max(V_i[s_{i+1}], R(s_i, a_i)) - K \\ &> K - K = 0. \end{aligned}$$

  So, there is a strict value increase, making the value of (s_i, a_i) (and thus the value of s_i) nonzero after the transition.

  The value of (s_i, a_i) will remain nonzero after all subsequent executions of (s_i, a_i). To see this, consider a transition

  $$(s_j, V_j) \xrightarrow{a_j,\, s_{j+1}} (s_{j+1}, V_{j+1}),$$

  where $j \geq i$ and $(s_j, a_j) = (s_i, a_i)$. Again, by Property 4.13 we have $\delta(s_i, a_i) \subseteq \mathit{strategy}(V_i) \subseteq \mathit{strategy}(V_j) = \mathit{strategy}(V)$. Also, we use Property D.5 and Property D.12 to know $V_j[s'] > K$ for each $s' \in \delta(s_i, a_i)$. Using Algorithm 1, where we substitute $V_j(s_i, a_i) = V_j[s_i]$ (since β selects only preferred actions), we have

  $$\begin{aligned} V_{j+1}(s_i, a_i) &= [\![V_j(s_i, a_i) + \max(V_j[s_{j+1}], R(s_i, a_i)) - K - V_j[s_i]]\!] \\ &= [\![\max(V_j[s_{j+1}], R(s_i, a_i)) - K]\!]. \end{aligned}$$

  Since $V_j[s_{j+1}] > K$, the right-hand side is strictly positive.

- Case 2.1 is not possible: the case would imply that $s_{i+1} \notin \mathit{strategy}(V_i) = \mathit{strategy}(V)$, which is false.

- Case 2.2.1 is not possible because $(s_i, a_i) \notin \mathit{rewards}(T)$.

- Case 2.2.2 is not possible, as we now explain. Towards a contradiction, suppose the case were applicable. The specific contradiction will be that although the first application of Case 2.2.2 leads to a change in the value of (s_i, a_i), any subsequent applications of (s_i, a_i) will keep the value fixed; so the above transition could occur only a finite number of times, which is false.

  To start, we note that the above transition changes the value of (s_i, a_i).

By Algorithm 1, we have:[20]

$$\begin{aligned} V_{i+1}(s_i, a_i) &= [\![V_i(s_i, a_i) + \max(V_i[s_{i+1}], R(s_i, a_i)) - K - V_i[s_i]]\!] \\ &= [\![\max(\mathit{expect}(V_i, s_i, a_i), 0) - K]\!] \\ &= [\![\mathit{expect}(V_i, s_i, a_i) - K]\!] \\ &= \mathit{expect}(V_i, s_i, a_i) - K. \end{aligned}$$

The last step uses $\mathit{expect}(V_i, s_i, a_i) - K > 0$ by Property D.5 (due to $S^{\text{start}} \not\subseteq \mathit{strategy}(V) = \mathit{strategy}(V_i)$). Case 2.2.2 implies $V_i(s_i, a_i) \neq \mathit{expect}(V_i, s_i, a_i) - K$. Hence, $V_i(s_i, a_i) \neq V_{i+1}(s_i, a_i)$.

We show that henceforth the value of (s_i, a_i) remains fixed. Suppose we encounter a subsequent transition t

$$(s_j, V_j) \xrightarrow{a_j,\, s_{j+1}} (s_{j+1}, V_{j+1}),$$

where $(s_j, a_j) = (s_i, a_i)$, and where still $V_j(s_i, a_i) = \mathit{expect}(V_i, s_i, a_i) - K$. We show concretely that $V_{j+1}(s_i, a_i) = V_j(s_i, a_i)$. We analyze why β has decided to perform transition t.

- Case 1 is not possible because $V_j[s_j] > 0$: indeed, $V_j[s_j] = V_j(s_i, a_i)$ because $a_j \in \mathit{pref}(s_j, V_j)$, and $V_j(s_i, a_i) = \mathit{expect}(V_i, s_i, a_i) - K > 0$ (as seen above).
- Case 2.1 is not possible: we have $\delta(s_i, a_i) \subseteq \mathit{strategy}(V_i)$ because we are working in Case 2.2.2, and $\mathit{strategy}(V_i) \subseteq \mathit{strategy}(V_j)$ (by Property 4.13).
- Case 2.2.1 is not possible since $(s_i, a_i) \notin \mathit{rewards}(T)$.
- Importantly, Case 2.2.2 is also not possible, as we now explain. Recall that $\delta(s_i, a_i) \subseteq \mathit{strategy}(V_i)$ because we are working in Case 2.2.2. Then Property D.12 (below) gives us $V_j[s'] = V_i[s']$ for each $s' \in \delta(s_i, a_i)$. Therefore $\mathit{expect}(V_j, s_i, a_i) = \mathit{expect}(V_i, s_i, a_i)$, and thus $V_j(s_i, a_i) = \mathit{expect}(V_j, s_i, a_i) - K$.
- Case 2.2.3.1 is not possible since $(s_i, a_i) \notin \mathit{rewards}(T)$.
- Only Case 2.2.3.2 is possible. We show that $V_{j+1}(s_i, a_i) = V_j(s_i, a_i)$. Inside Algorithm 1, we have

  $$d = \max(V_j[s_{j+1}], R(s_i, a_i)) - K - V_j[s_j].$$

  It suffices to show $d = 0$. In the equation for d, the following substitutions can be done:

  1. $V_j[s_j] = V_j(s_i, a_i) = \mathit{expect}(V_i, s_i, a_i) - K$ since $a_j = a_i \in \mathit{pref}(s_j, V_j)$;
  2. $V_j[s_{j+1}] = \mathit{expect}(V_j, s_i, a_i) = \mathit{expect}(V_i, s_i, a_i)$ by Case 2.2.3.2;[21]
  3. $R(s_i, a_i) = 0$.

[20] We substitute (1) $V_i(s_i, a_i) = V_i[s_i]$ since $a_i \in \mathit{pref}(s_i, V_i)$, and (2) $V_i[s_{i+1}] = \mathit{expect}(V_i, s_i, a_i)$ by design of Case 2.2.2.

[21] The equality $\mathit{expect}(V_j, s_i, a_i) = \mathit{expect}(V_i, s_i, a_i)$ can be seen with the same reasoning as in the discussion of Case 2.2.2 just above.

Hence,

$$\begin{aligned} d &= \mathit{expect}(V_i, s_i, a_i) - K - (\mathit{expect}(V_i, s_i, a_i) - K) \\ &= 0. \end{aligned}$$

- Case 2.2.3.1 is not possible because $(s_i, a_i) \notin \mathit{rewards}(T)$.
- Case 2.2.3.2 is possible. We show $s_i \in \mathit{strategy}(V_i)$, which, combined with $s_i \notin \mathit{strategy}(V)$, gives the desired contradiction. Concretely, we show $s_i \in z_{j+1}(V_i)$ for some $j \geq 1$, which implies $s_i \in \mathit{strategy}(V_i)$.

  We first show $V_i[s_i] = M - (j+1)K$ for some $j \geq 1$. Recall that $(s_i, a_i) \notin \mathit{rewards}(T)$. We have $\delta(s_i, a_i) \subseteq \mathit{strategy}(V_i)$ because Case 2.1 was not applicable. Since Case 2.2.2 was not applicable, we know

  $$V_i(s_i, a_i) = \mathit{expect}(V_i, s_i, a_i) - K.$$

  Let $s' \in \delta(s_i, a_i)$ with $V_i[s'] = \mathit{expect}(V_i, s_i, a_i)$. Since $s' \in \mathit{strategy}(V_i)$, we can consider the smallest index j satisfying $s' \in z_j(V_i)$; note that $j \geq 1$. Hence, $V_i[s'] = M - jK$. We note the following, where we start with $V_i[s_i] = V_i(s_i, a_i)$ since β only chooses preferred actions:

  $$\begin{aligned} V_i[s_i] &= V_i(s_i, a_i) \\ &= \mathit{expect}(V_i, s_i, a_i) - K \\ &= V_i[s'] - K \\ &= M - jK - K \\ &= M - (j+1)K. \end{aligned}$$

  Next, we show that the actions in $\mathit{pref}(s_i, V_i)$ satisfy the desired properties, in the definition of $z_{j+1}(V_i)$ in Section 4.5.2. Let $a \in \mathit{pref}(s_i, V_i)$.

  1. Since $V_i[s_i] = M - (j+1)K$ and $j \geq 1$, we have $V_i[s_i] < M - K$. We have $(s_i, a) \notin \mathit{rewards}(T)$: otherwise, because Case 2.2.1 was not applicable (as mentioned above), we would have $V_i[s_i] \geq V_i(s_i, a) = M - K$, which is false.
  2. Because Case 2.1 was not applicable, and $(s_i, a) \notin \mathit{rewards}(T)$ (see the previous item), we know $\delta(s_i, a) \subseteq \mathit{strategy}(V_i)$.
     Since $a \in \mathit{pref}(s_i, V_i)$, we have

     $$V_i(s_i, a) = V_i[s_i] = M - (j+1)K.$$

     Moreover, since Case 2.2.2 was not applicable, we know

     $$V_i(s_i, a) = \mathit{expect}(V_i, s_i, a) - K.$$

     By combining the above two expressions for $V_i(s_i, a)$, we know

     $$\mathit{expect}(V_i, s_i, a) = M - jK.$$

     Since $\delta(s_i, a) \subseteq \mathit{strategy}(V_i)$, we therefore know $\delta(s_i, a) \subseteq z_j(V_i)$.[22]

[22] Suppose there is some $s'' \in \delta(s_i, a) \setminus z_j(V_i)$. Then there is some smallest index k with $s'' \in z_k(V_i)$ where $k > j$. Then $V_i[s''] = M - kK < M - jK$, implying $\mathit{expect}(V_i, s_i, a) < M - jK$, which is false.

3. Since $\mathit{expect}(V_i, s_i, a) = M - jK$ (see previous item), there must be some state $s'' \in \delta(s_i, a)$ with $V_i[s''] = M - jK$.

Property D.12. Function β preserves the value of strategy states. More formally, consider a sequence of transitions generated by function β,

$$(s_1, V_1) \xrightarrow{a_1,\, s_2} \dots \xrightarrow{a_{n-1},\, s_n} (s_n, V_n).$$

For each $s \in \mathit{strategy}(V_1)$ we have $V_1[s] = V_n[s]$.

Proof. For each transition

$$(s_i, V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1}, V_{i+1}),$$

with $i \in \{1, \dots, n-1\}$, we know the following:

- By Property 4.13, we know $\mathit{strategy}(V_i) \subseteq \mathit{strategy}(V_{i+1})$.
- By Property D.11, we know $V_{i+1}(s, a) = V_i(s, a)$ for each $(s, a) \in \mathit{strategy}(V_i) \times A$.

Now, we fix some $s \in \mathit{strategy}(V_1)$. We have

- $s \in \mathit{strategy}(V_1) \subseteq \mathit{strategy}(V_2)$ and $V_1[s] = V_2[s]$;
- $s \in \mathit{strategy}(V_2) \subseteq \mathit{strategy}(V_3)$ and $V_2[s] = V_3[s]$;
- $\dots$

By transitivity, $V_1[s] = V_n[s]$, as desired. □

D.7 Proof of Property 4.13

Consider a transition generated by β,

$$(s_i, V_i) \xrightarrow{a_i,\, s_{i+1}} (s_{i+1}, V_{i+1}).$$

By Property D.11 we know the following: $\forall s \in \mathit{strategy}(V_i)$, $\forall a \in A$,

$$V_{i+1}(s, a) = V_i(s, a).$$

This implies for all $s \in \mathit{strategy}(V_i)$ that

1. $V_{i+1}[s] = V_i[s]$; and,
2. $\mathit{pref}(s, V_{i+1}) = \mathit{pref}(s, V_i)$.

We now show by induction on $j = 1, 2, \dots$ that $z_j(V_i) \subseteq z_j(V_{i+1})$, resulting in $\mathit{strategy}(V_i) \subseteq \mathit{strategy}(V_{i+1})$, as desired.

- For the base case, let $s \in z_1(V_i)$. This implies (1) $V_{i+1}[s] = V_i[s] = M - K$, and (2) for each $a \in \mathit{pref}(s, V_{i+1}) = \mathit{pref}(s, V_i)$ that $(s, a) \in \mathit{rewards}(T)$. Hence, $s \in z_1(V_{i+1})$.
- Let $j \geq 2$. The induction hypothesis is $z_{j-1}(V_i) \subseteq z_{j-1}(V_{i+1})$. For the inductive step, let $s \in z_j(V_i) \setminus z_{j-1}(V_i)$. We have (1) $V_{i+1}[s] = V_i[s] = M - jK$ and (2) for all $a \in \mathit{pref}(s, V_{i+1}) = \mathit{pref}(s, V_i)$,

1. $(s, a) \notin \mathit{rewards}(T)$;
2. $\delta(s, a) \subseteq z_{j-1}(V_i)$, which, combined with the induction hypothesis $z_{j-1}(V_i) \subseteq z_{j-1}(V_{i+1})$, gives $\delta(s, a) \subseteq z_{j-1}(V_{i+1})$;
3. $\exists s' \in \delta(s, a)$ with $V_{i+1}[s'] = V_i[s'] = M - (j-1)K$.[23]

Overall, we see $s \in z_j(V_{i+1})$.

D.8 Proof of Property 4.14

Suppose we perform β infinitely often, starting at some arbitrary configuration. Towards a contradiction, suppose that after a while we no longer encounter transitions with reward.

Because there are only a finite number of configurations in any infinite transition sequence (Lemma 2.7), we encounter a configuration-cycle $\mathcal{C}$,

$$(s_1, V_1) \xrightarrow{a_1,\, s_2} \dots \xrightarrow{a_{n-1},\, s_n} (s_n, V_n),$$

where $n \geq 2$, and $(s_1, V_1) = (s_n, V_n)$, and where none of the transitions contains reward.

Property D.13 (below) tells us that $V_n[s_n] = 0$. Now, because β is deterministic, and $(s_1, V_1) = (s_n, V_n)$, the cycle $\mathcal{C}$ gives rise to another cycle $\mathcal{C}'$ that is shifted one transition into the future:

$$(s_2, V_2) \xrightarrow{a_2,\, s_3} \dots \xrightarrow{a_{n-1},\, s_n} (s_n, V_n) \xrightarrow{a_n,\, s_{n+1}} (s_{n+1}, V_{n+1}),$$

where (s_2, V_2) is the second configuration in $\mathcal{C}$, $(a_n, s_{n+1}) = (a_1, s_2)$, and $(s_2, V_2) = (s_{n+1}, V_{n+1})$ (because $(s_1, V_1) = (s_n, V_n)$). All state-action pairs that are executed in cycle $\mathcal{C}'$ are also executed in cycle $\mathcal{C}$; hence, $\mathcal{C}'$ contains no reward either. Property D.13 again gives $V_{n+1}[s_{n+1}] = 0$. The reasoning can now be repeated arbitrarily many times, to establish an infinite (and contiguous) sequence of configurations in which the current state has zero value.[24] This means that all transitions are specifically generated by Case 1 of function β. Intuitively, we have designed β in such a way that if the current state has value zero then we move the agent towards reward. We look at the sub-cases of Case 1:

- Case 1.1: in that case we follow the acyclic movement strategy defined for the non-reducible states. Eventually we encounter a start state, which is reducible by assumption. We therefore must eventually arrive at one of the following two cases.
- Case 1.2. In that case we obtain reward, which would be a contradiction.
- Case 1.3. In that case we move strictly deeper into the reducibility layers. However, this process can not continue forever because there are only a finite number of states. We must eventually arrive at Case 1.2 and obtain reward (again, a contradiction).

In this case analysis, we have therefore arrived at the desired contradiction.

[23] Here we use $s' \in z_{j-1}(V_i) \subseteq \mathit{strategy}(V_i)$, which implies $V_{i+1}[s'] = V_i[s']$.

[24] If $n = 2$ then all configurations after (s_1, V_1) are also (s_1, V_1); the current state therefore has value zero forever.

Property D.13. Suppose we have configuration-cycle $\mathcal{C}$ under function β,

$$(s_1, V_1) \xrightarrow{a_1,\, s_2} \dots \xrightarrow{a_{n-1},\, s_n} (s_n, V_n),$$

where $(s_1, V_1) = (s_n, V_n)$. If none of the transitions contains reward then $V_n[s_n] = 0$. Intuitively, this means that value can not be sustained in absence of reward.

Proof. Assume that none of the transitions contains reward. Abbreviate $P = \{s_1, \dots, s_{n-1}\}$. We consider the highest value in V_1 among the states in P:

$$W = \max \{V_1[s] \mid s \in P\}.$$

We show $W = 0$. This implies in particular that $V_1[s_1] = V_n[s_n] = 0$, as desired.

Assume for now that we already know the following:

Claim D.14. For each $i \in \{1, \dots, n-1\}$, we have $V_{i+1}(s_i, a_i) < W$.

Now, let $i \in \{1, \dots, n-1\}$ be the smallest index for which $V_1[s_i] = W$, i.e., index i is the first index at which we encounter a state with value W in V_1. Because s_i could not have been encountered before in the cycle (by choice of i), we have $V_i[s_i] = V_1[s_i] = W$. Moreover, because function β always chooses a preferred action, we have $a_i \in \mathit{pref}(s_i, V_i)$, which implies $V_i(s_i, a_i) = V_i[s_i]$. Hence, $V_i(s_i, a_i) = W$. Again, because s_i could not have been visited before (by choice of i), we have $V_1(s_i, a_i) = W$. By Claim D.14, we know $V_{i+1}(s_i, a_i) < W$. All subsequent transitions for the pair (s_i, a_i) result in a value strictly smaller than W. This results in $V_n(s_i, a_i) < W$. But since $V_1(s_i, a_i) = W$, we would have $V_n \neq V_1$; this is the desired contradiction.

Proof of Claim D.14. We first consider the following sub-claim:

Claim D.15. Let $i \in \{1, \dots, n-1\}$. Assume for each $s \in P$ that $V_i[s] \leq W$. Then $V_{i+1}(s_i, a_i) < W$.

To finish the proof of Claim D.14, we show by induction on $i \in \{1, \dots, n-1\}$ that, for each $s \in P$, we have $V_i[s] \leq W$. For the base case, for each $s \in P$, we have $V_1[s] \leq W$ by definition of W. For the inductive step, let $i \geq 1$, with the assumption $V_i[s] \leq W$ for each $s \in P$. Letting $s \in P$, we observe that the desired property is satisfied for V_{i+1}:

- If $s \neq s_i$ then $V_{i+1}[s] = V_i[s] \leq W$ by the induction hypothesis.
- Suppose $s = s_i$. For any $a \in A$ if $a \neq a_i$ then $V_{i+1}(s, a) = V_i(s, a) \leq W$ by the induction hypothesis; if $a = a_i$ then $V_{i+1}(s, a) < W$ by Claim D.15 (using also the induction hypothesis).

To show Claim D.15, let $i \in \{1, \dots, n-1\}$. By Algorithm 1, we have

$$V_{i+1}(s_i, a_i) = [\![V_i(s_i, a_i) + \max(V_i[s_{i+1}], R(s_i, a_i)) - K - V_i[s_i]]\!].$$

We have $V_i(s_i, a_i) = V_i[s_i]$ since function β always chooses an action that is preferred by the agent. Also, we have $R(s_i, a_i) = 0$ because none of the transitions contains reward by assumption. Moreover, always $V_i[s_{i+1}] \geq 0$. The equation can now be simplified as follows:

$$V_{i+1}(s_i, a_i) = [\![V_i[s_{i+1}] - K]\!].$$

Since $s_{i+1} \in P$, and $V_i[s] \leq W$ for each $s \in P$ (by the assumption in Claim D.15), we obtain

$$\begin{aligned} V_{i+1}(s_i, a_i) &\leq [\![W - K]\!] \\ &< W. \end{aligned}$$

In the last step, we use $W - K \leq W$, which implies $[\![W - K]\!] \leq [\![W]\!] = W$ (using that $W \geq 0$). □